BBC RD 1985/3

RESEARCH DEPARTMENT

REPORT

Further developments of CCD line array telecine

I. Childs, B.A., Ph.D., M.I.E.E., C.Eng.

Research Department, Engineering Division
THE BRITISH BROADCASTING CORPORATION

February, 1985

BBC RD 1985/3
UDC 621.377.6:
621.385.832.5

FURTHER DEVELOPMENTS OF CCD LINE-ARRAY TELECINE

I. Childs, B.A., Ph.D., MIEE, C.Eng.

Summary

In recent years there have been several advances in technology relevant to the use of solid-state line-array sensors in telecine applications. This Report will examine the improvements in performance of recent generations of line-array sensors, together with the further benefits obtainable by using improved signal processing. The use of digital storage as a necessary part of a line-array telecine also opens up additional capabilities for facilities such as variable-speed running and the concealment of dirt and scratches; both of these features are considered in this Report.

Issued under the Authority of

Head of Research Department

Research Department, Engineering Division,
BRITISH BROADCASTING CORPORATION
February 1985
(PH-254)

This Report may not be reproduced in any form without the written permission of the British Broadcasting Corporation.

It uses SI units in accordance with B.S. document PD 5686.

FURTHER DEVELOPMENTS OF CCD LINE-ARRAY TELECINE

I. Childs, B.A., Ph.D., M.I.E.E., C.Eng.

Section	Title	Page
	Summary	Title Page
1.	Introduction	1
2.	Factors affecting the picture quality of line-array telecines	1
	2.1. Improved sensors	1
	2.1.1. Operation of charge-coupled devices	1
	2.1.2. Characteristics of CCPD sensors	3
	2.2. The use of correlated double sampling to enhance signal-to-noise ratio	4
	2.3. Improved techniques for concealing sensitivity variations	6
3.	Variable speed operation	8
	3.1. Basic control system	8
	3.2. Compensation for integration time variations	10
	3.3. Operation at slow film speeds	11
	3.4. Operation at fast film speeds	12
	3.5. Movement interpolation	15
4.	Dirt concealment	15
	4.1. Detection of dirt by infra-red light	16
	4.2. Concealment strategies	18
	4.2.1. Variable-gain approach	18
	4.2.2. Concealment by substitution	19
	4.3. Experimental equipment	19
5.	Conclusions	24
6.	Acknowledgements	24
7.	References	24

© BBC 2006. All rights reserved. Except as provided below, no part of this document may be reproduced in any material form (including photocopying or storing it in any medium by electronic means) without the prior written permission of BBC Research & Development except in accordance with the provisions of the (UK) Copyright, Designs and Patents Act 1988.

The BBC grants permission to individuals and organisations to make copies of the entire document (including this copyright notice) for their own internal use. No copies of this document may be published, distributed or made available to third parties whether by paper, electronic or other means without the BBC's prior written permission. Where necessary, third parties should be directed to the relevant page on BBC's website at http://www.bbc.co.uk/rd/pubs/ for a copy of this document.

FURTHER DEVELOPMENTS OF CCD LINE-ARRAY TELECINE

I. Childs, B.A., Ph.D., MIEE, C.Eng.

1. Introduction

It was in 1973 that the use of line-array solid state sensors as image pick-up devices in telecine machines was first suggested;[1] experimental results were reported which had been gained from the use of a 512-element photodiode array. Improvements to the signal processing associated with the use of photodiode sensors were investigated[2] but performance was limited by a fundamental drawback associated with photodiode devices; the driving waveforms have large in-band frequency components and any spurious breakthrough into the video output is extremely difficult to remove.

By 1976, however, early examples of sensors using a different principle had been reported;[3] these were charge-coupled devices (CCD's) and had the advantage that the driving waveforms had no in-band frequency components and thus spurious breakthrough on the video output could be removed by a simple lowpass filter. These first CCD's were surface channel devices in which the electrons making up the output signal were transported along the surface of the silicon. Later CCD's used the buried channel approach, whereby the signal electrons were transported just below the silicon surface;[4] the increased charge transfer efficiency of this technique yielded a greatly improved response at high video frequencies. By 1978, therefore, most of the problems associated with the use of CCD line-arrays in a monochrome telecine had been overcome.

The next stage was the use of three such CCD's to realise a colour telecine;[5] each sensor had 1024 elements, more than enough to match the 5.5 MHz video bandwidth required. The only problems were found to be the poor spectral response of the sensors (which meant that they had to be selected for the red, green and blue channels) and their relatively high level of electrical noise (particularly in view of their poor blue sensitivity).

In the last few years, however, there have been several significant advances, both in the technology used to manufacture sensors and in the signal processing of the telecine; these advances will be described in this Report. The first section will cover improvements affecting the basic picture quality obtainable; of particular interest are the characteristics of recent photodiode/CCD hybrid sensors and the use of noise cancelling circuits at the sensor output. The second section will describe the way in which line-array sensors may be coupled with the other parts of the telecine in order to yield a system capable of operation over a wide range of film speeds and film gauges. The final section will describe the use of the good infra-red sensitivity of sensors to detect the presence of physical blemishes, such as dirt and scratches. Once detected, these blemishes may be concealed by substituting the corrupted video signals with signals generated according to a suitable strategy, and some suggestions will be made on possible choices for such a strategy.

2. Factors affecting the picture quality of line-array telecines

2.1. Improved sensors

2.1.1. Operation of charge-coupled devices

Although the detailed physics of the construction and operation of charge-coupled devices is outside the scope of this report, a brief description will be given in order to enable the reader who is unfamiliar with this subject to understand some of the later sections. For further information the reader is referred to some of the excellent books on the subject.[6]

Fig. 1 (a) shows the construction of a three-phase charge-coupled device. It consists of a p- or n-type channel in the silicon substrate; covering this channel, but insulated from it by a thin layer of silicon oxide, are three sets of electrodes. If the appropriate voltage is applied to one set of electrodes a series of potential wells is formed in the channel beneath them and minority carriers injected into these wells are unable to leave and are thus stored.

If the positions of the wells are altered by varying the voltages on the three sets of electrodes as in Figs. 1 (b) and 1 (c), the stored charges can be moved along the channel. The CCD thus behaves as an analogue shift register, so that charges in each of the wells may be moved along the register to one end, at which point they may be fed to an output circuit.

The charges may enter the wells either from an input gate or as carriers released when light is allowed to fall on the substrate. It is this latter effect

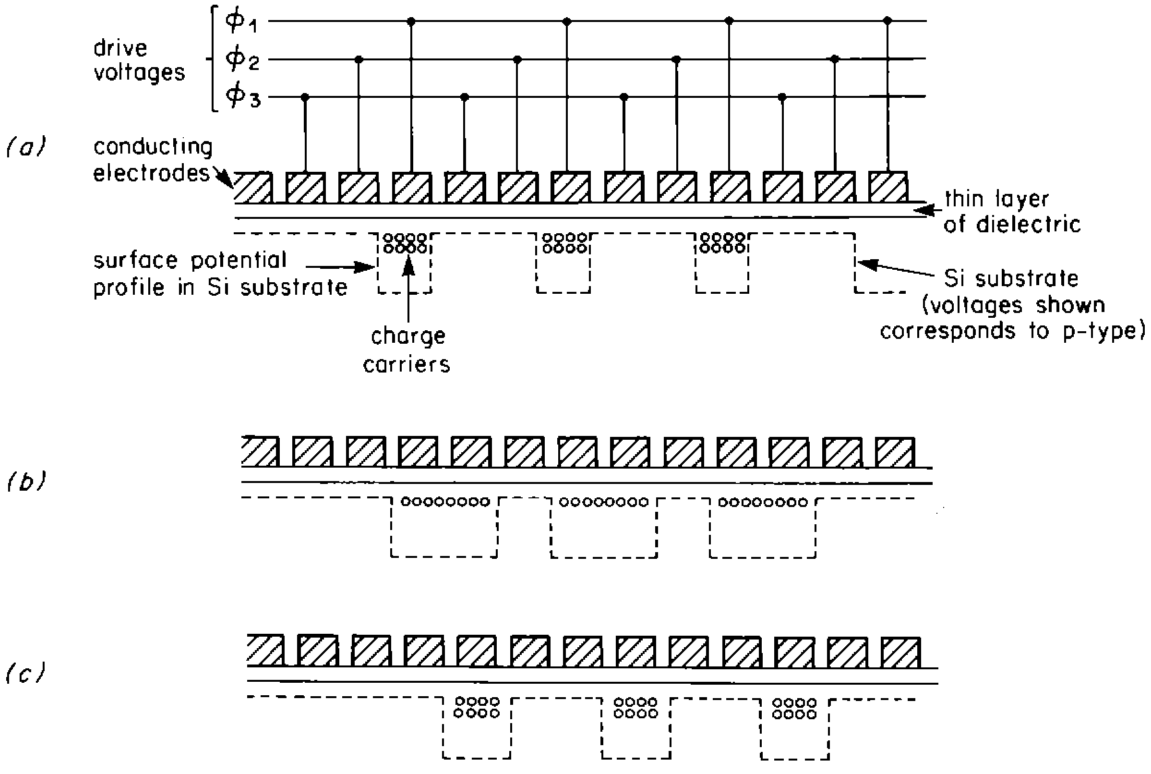

Fig. 1 – Construction and operation of 3-phase CCD (a) ϕ_1 voltage positive, ϕ_2 and ϕ_3 voltages zero, (b) ϕ_2 taken positive ϕ_1 and ϕ_3 unchanged, (c) ϕ_2 remains positive, ϕ_1 taken to zero.

which is used in image-forming applications. Fig. 2 shows an arrangement commonly used in single-line arrays. Light falling on the image sensing elements releases charge which is stored in potential wells underneath each element. Once per television line the contents of all the individual elements are simultaneously passed to corresponding elements in a transfer register. While the next charge image is being formed under the image sensing elements, the contents of the transfer register are read out in sequence, and form an electrical signal. The two parts of the CCD thus perform the image sensing

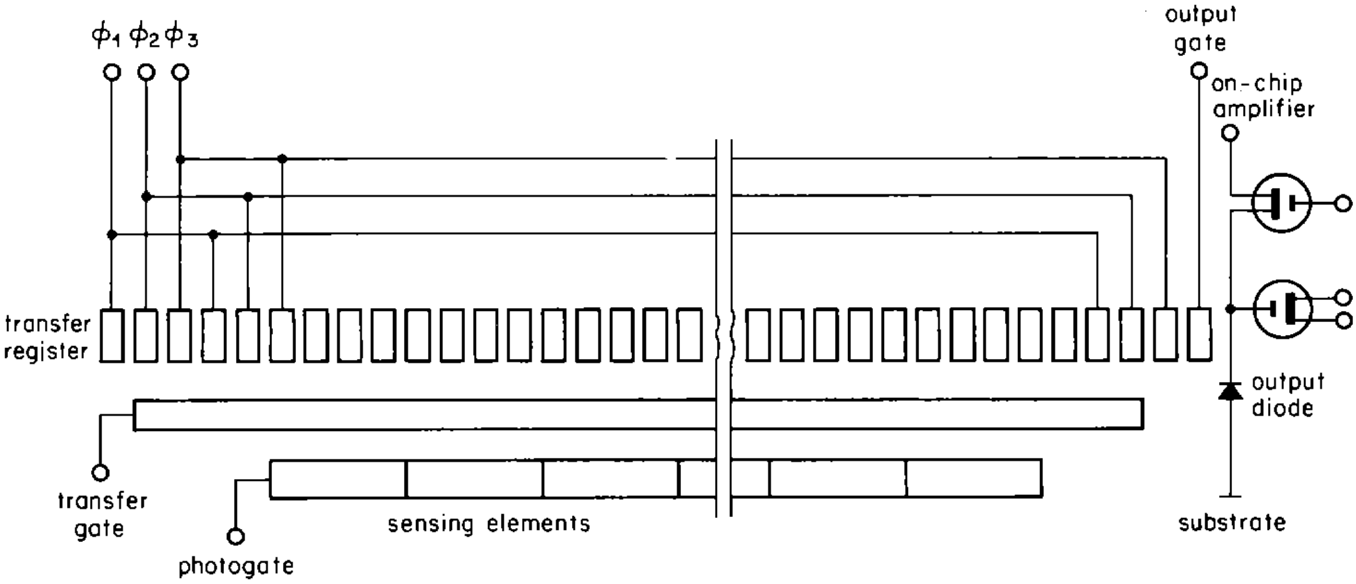

Fig. 2 – Layout of a typical linear sensor array.

and one of the scanning functions required to produce a television signal.

There are several variations possible. If the profile of the potential wells in the channel is made asymmetric, either by implanting impurities or by varying the thickness of the oxide layer, it is possible to achieve charge transport with only two clock phases instead of three.[7,8] This considerably simplifies the electrode structure by removing the need for one set of electrodes to cross over another.

A second variation is the buried channel device.[9,10] In a surface-channel CCD, defects in the silicon trap some of the carriers so that not all the charge is transferred. The concentration of these defects is higher at the surface than in the bulk material. Thus, if the channel is formed below the surface of the substrate by additional diffusions during manufacture, the proportion of charge transferred at each step is improved. An additional benefit obtained by increasing the spacing between the electrode structure and the channel is that the profile of the electric field from each electrode in the charge-carrying region is made more suitable for high-speed operation.

From this description, it is clear that the light must pass through the electrodes on the surface of the CCD before it is detected. This is usually achieved by making them of some transparent material such as polysilicon. However, the use of polysilicon gives rise to two disadvantages. The first is that its refractive index is not a good match between that of air and that of the silicon oxide beneath; thus an interference filter is produced which creates irregular ripples in the spectral response. The second disadvantage is that polysilicon has a poor transmission to blue light; thus the sensitivity of CCD line-arrays to blue wavelengths is poor. Both these features of CCD sensors using polysilicon photogates are shown in the spectral response curve of Fig. 3 (a), taken from reference 5.

Several techniques have been suggested in order to combat these problems. To date, the most successful of these is the use of photodiode sensing elements (which have no covering polysilicon layer) coupled with a CCD transfer register. Such a hybrid device is usually referred to as a CCPD (charge-coupled photodiode) sensor.

2.1.2. Characteristics of CCPD sensors

The spectral response characteristics of two CCPD devices, made by different manufacturers were measured; the results are shown in Figs 3 (b) and 3 (c). These characteristics revealed differences in the fabrication process, since one CCPD (sensor 1) still retained ripples in its response while these were largely eliminated in the second device (sensor 2). Both CCPD's showed a significant improvement

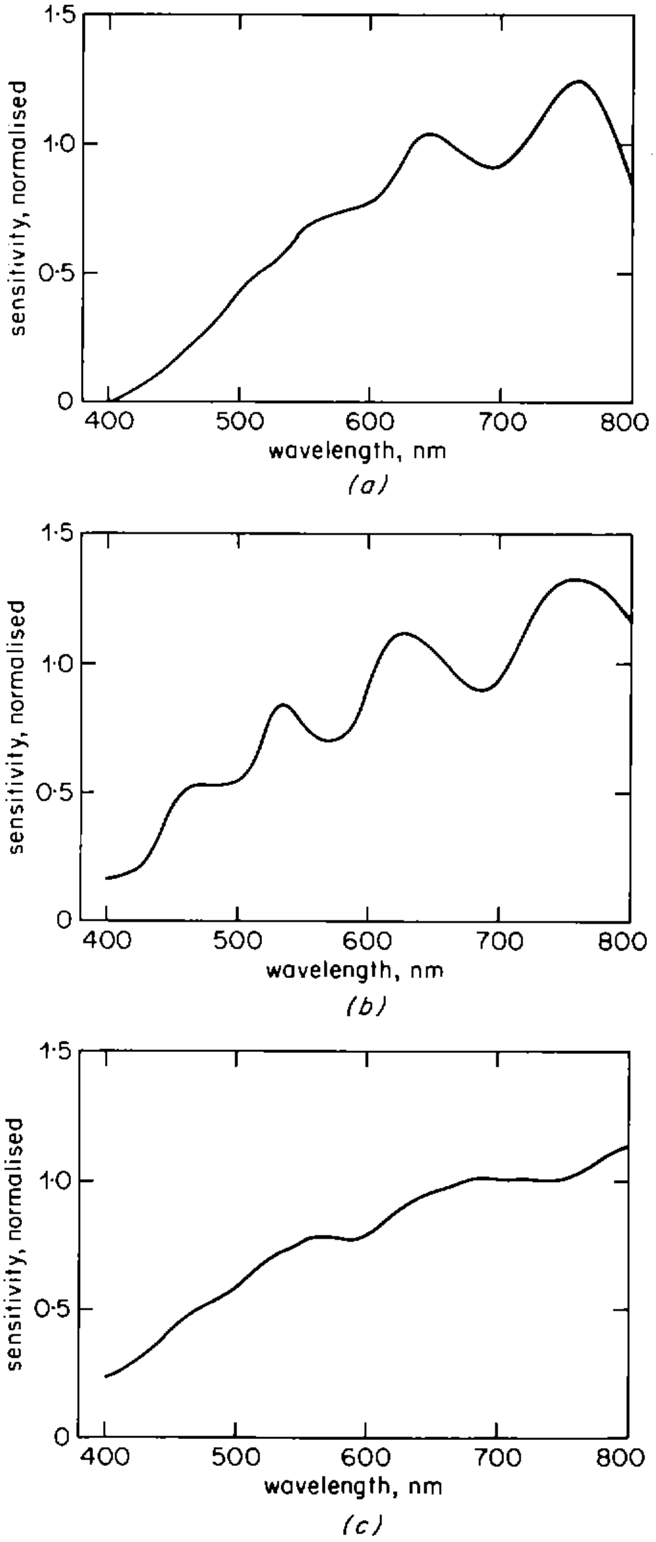

Fig. 3 – Spectral sensitivity of three linear arrays, (a) CCD sensor, (b) CCPD sensor (sensor 1), (c) CCPD sensor made by a different manufacturer (sensor 2).

in blue response over the earlier CCD, however; in the case of sensor 2, this was coupled with an overall increase in dynamic range due to the use of a redesigned output stage. Such an improved blue sensitivity can be used not only to increase the overall signal-to-noise ratio, but also to improve the colorimetry of the telecine and reduce any unwanted response due to infra-red light.

One problem was, however, found to occur in the CCPD sensors. This was known as photosite transfer loss and occurred at the interface between the photodiode sensing elements and the CCD readout register. Instead of the charge collected in each sensing element being totally transferred to the readout register, some was left behind; it was then read out on subsequent scans. This effect was measured and the results for the CCD and one of the CCPD devices are shown in Fig. 4(a), (b), (c) and (d); a light-emitting diode illuminated the line-array sensor with a short pulse of light lasting for less than a television line. It can be seen that, although the CCD sensor had negligible output on subsequent lines, the CCPD sensor did show a measurable, but small, signal. This spurious signal tended to be of a fixed amplitude, independent of the size of the original main pulse. Fig. 5 shows the output of the CCPD sensor in response to a pulse of only 10% of that used for the results of Fig. 4(c) and (d); the amplitude of the spurious response was only slightly different from that previously observed. The effect of this photosite transfer loss on the picture was to produce vertical smearing, or loss of detail, which was most obvious in dark areas. In extreme cases it could even appear as if the sensor had no response to very dark, low-level detail signals since these were swamped by the spurious response from previous lines. To date no method for overcoming photosite transfer loss in CCPD sensors has been reported, other than by the use of bias light. Because of the problems associated with the incorporation of bias light, however, it is usually preferable to accept a small amount of transfer loss which can be minimised by careful attention to drive waveforms.

In most other respects there is little to distinguish the performance of a CCPD from that of a CCD device. For this reason, in the remainder of this Report the term "CCD" will be used to encompass both types of sensor unless specifically indicated.

2.2. The use of correlated double sampling to enhance signal-to-noise ratio

There are many sources of noise in a CCD sensor; some of the most important of these are indicated in Table 1. Two, however, are predominant. The first is photon noise arising from the discrete nature of the incident illumination. It has a flat frequency spectrum and its amplitude is related to the square root of the instantaneous light level. In reference (5) it was estimated that the maximum noise produced by this mechanism was approximately −58 dB with respect to the signal output of the sensor when the sensor was illuminated to near saturation level. The improvement in detection efficiency gained by the newer CCPD devices can be expected to improve the signal-to-noise ratio by one or two dB because the incident photons are no longer being reflected by, or absorbed in, the polysilicon electrode layer. This level of photon noise is imperceptible under normal circumstances.

Far more disturbing is the electrical noise generated in the output stage of the CCD. Fig. 6(a) shows a functional circuit diagram of one of the most usual types of output amplifier. The size of the charge packet in each potential well is determined by the extent to which it discharges an on-chip capacitor (a reverse-biased diode junction). After each charge packet has been detected the potential on the diode is reset by a MOSFET switch prior to the detection of the following charge packet. Noise is added both by the switch resistance of the MOSFET and by the source follower stages used to buffer the diode output. The source follower noise, in particular, is extremely disturbing since it tends to have a predominant low-frequency content giving the noise the unpleasant "streaky" appearance shown in Fig. 7(a)*. As has been noted in reference 5 this low-frequency noise is much more disturbing than its low level would suggest.

Fortunately, there is a technique that can be used to reduce the objectionable low frequency noise components. Fig. 6(b) shows the output waveform from the circuit of Fig. 6(a) and it is clear that the low-frequency noise components will be highly correlated between instants (a) and (b). Thus, if the signal is sampled at these two instants and the samples are subtracted, the noise will be cancelled but the signal will be unaffected. This technique is known as correlated double sampling.[6]

The effect of using this technique on the appearance of the noise is shown in Fig. 7(b); it is subjectively much less visible than in Fig. 7(a). Fig. 8 shows the frequency spectrum of the noise with, and without, correlated double sampling and confirms that the excess low-frequency components have been removed; the noise spectrum after correlated double sampling is substantially flat. This white noise is much more easily masked by film grain than

---

*N.B. In order to increase the visibility of the noise in Figs. 7(a) and 7(b), the signal was amplified by a further 20 dB after gamma correction.

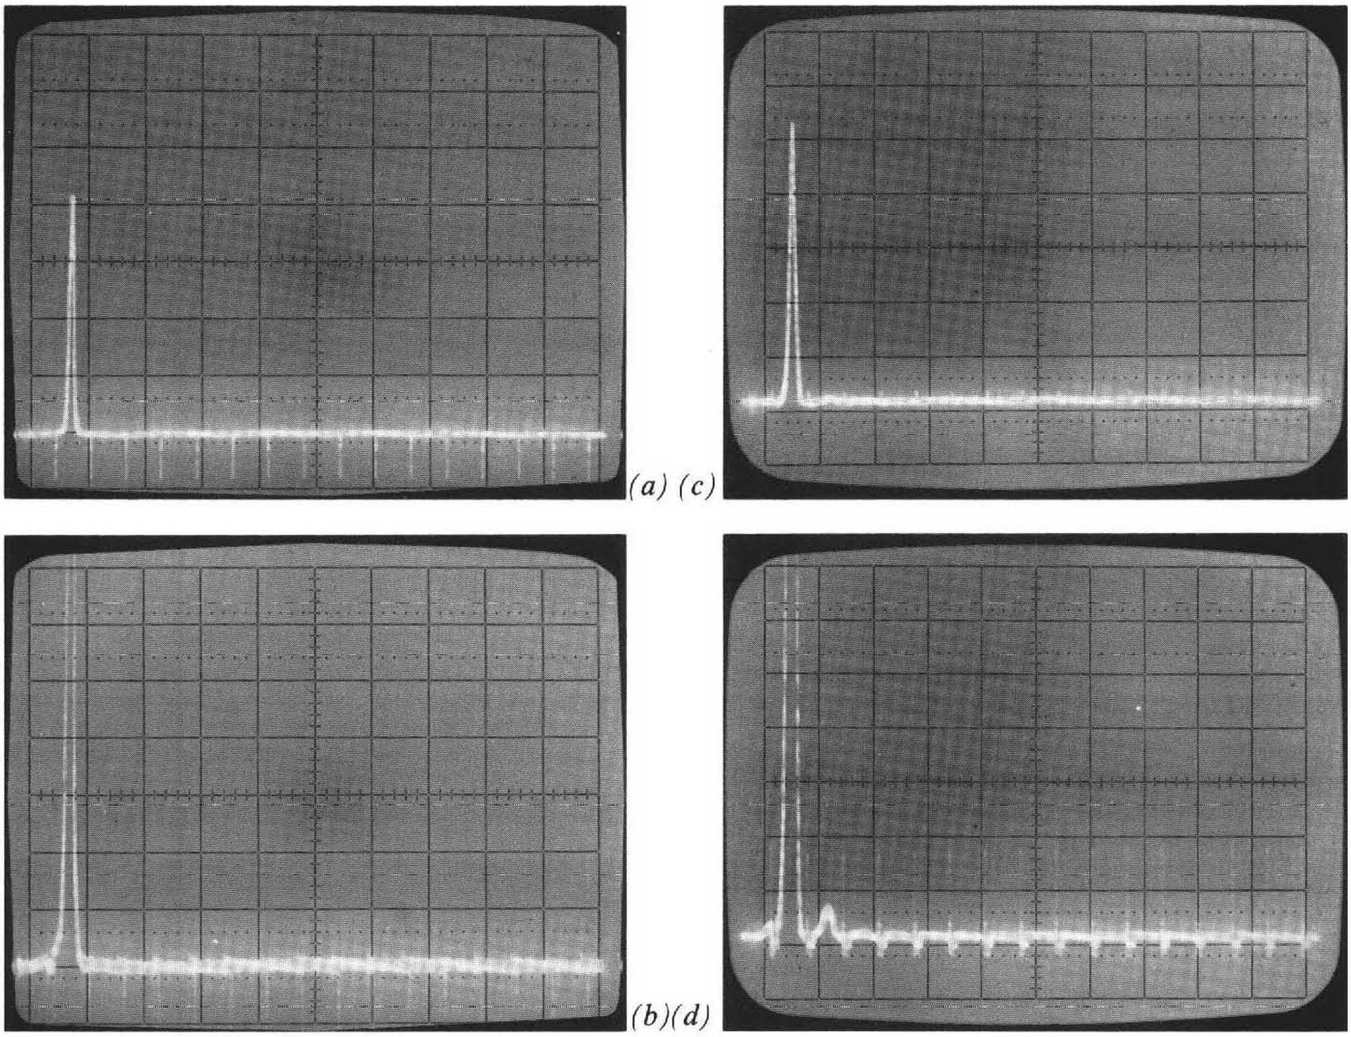

(a) (c)

(b)(d)

Fig. 4 – Photosite transfer loss of linear arrays, (a) CCD sensor. Scales: hor. 100 μsec/div, vert. 20% V_{sat}/div, (b) as (a) but vertical scale × 10 (2% V_{sat}/div), (c) CCPD sensor. Scales: hor. 100 μsec/div, vert 20% V_{sat}/div, (d) as (c) but vertical scale × 10 (2% V_{sat}/div).

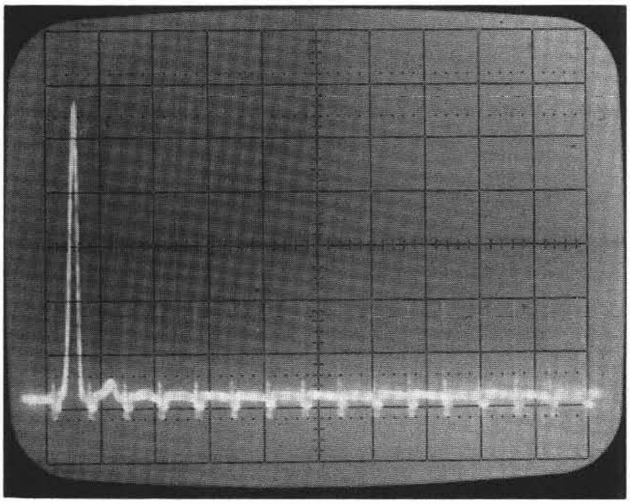

Fig. 5 – Photosite transfer loss of CCPD sensor at low light levels. Scales: as Figs. 4(b) and (d) (hor. 100 μsec/div, vert. 2% V_{sat}/div).

TABLE 1
SOURCES OF NOISE IN A CCD SENSOR

---

1. Input noise – photon noise
   – noise due to any electrical bias ("fat zero")
2. Transfer noise
3. Dark current noise
4. Noise in output amplifier – reset noise
   – noise in MOSFET buffer

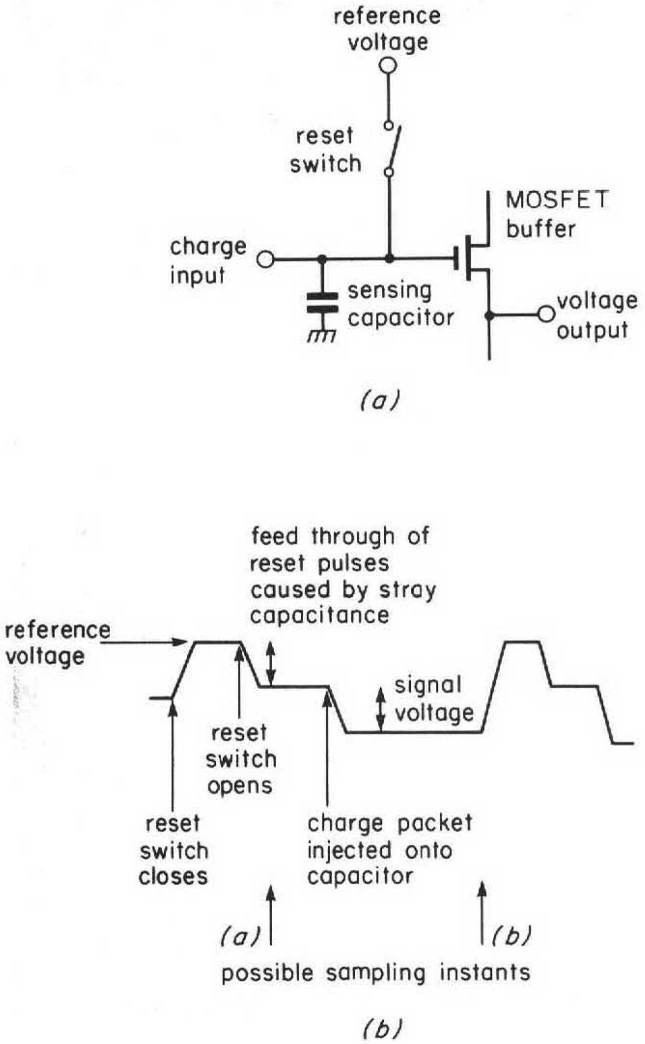

Fig. 6 – Output amplifier of CCD, (a) typical circuit, (b) output voltage waveform.

the streaky noise at the sensor output; thus the use of correlated double sampling substantially increases the useful dynamic range of the telecine.

2.3. Improved techniques for concealing sensitivity variations

Because of mask tolerances in the manufacturing process and variations in the silicon substrate, the sensitivities of the individual photosites are not all identical. In addition any vignetting or optical shading in the telecine optical system results in non-uniform illumination of the sensors. The effects of all these factors is to modulate the picture information with a fixed pattern of vertical stripes; the appearance of these stripes, and the techniques for their concealment, has already been covered in reference 5. In essence, the signal output from each element must be amplified by a variable amount, the gain applied to each element being adjusted to compensate for the particular sensitivity of that element. Because of its action in reducing the visibility of vertical stripes the compensation circuit was termed a "stripe-stripper".

The system described in reference 5 (and shown in Fig. 9) suffered from a number of drawbacks. In the first instance, it processed a logarithmically-converted signal. Such a signal was most accurately generated digitally but required an analogue-to-digital converter with up to 11 bits resolution;[11] the design of such a converter presents a considerable challenge. Use of analogue circuitry to implement the necessary logarithmic converters was found to give inadequate performance; the inaccuracy of the transfer characteristic in dark areas caused a direct

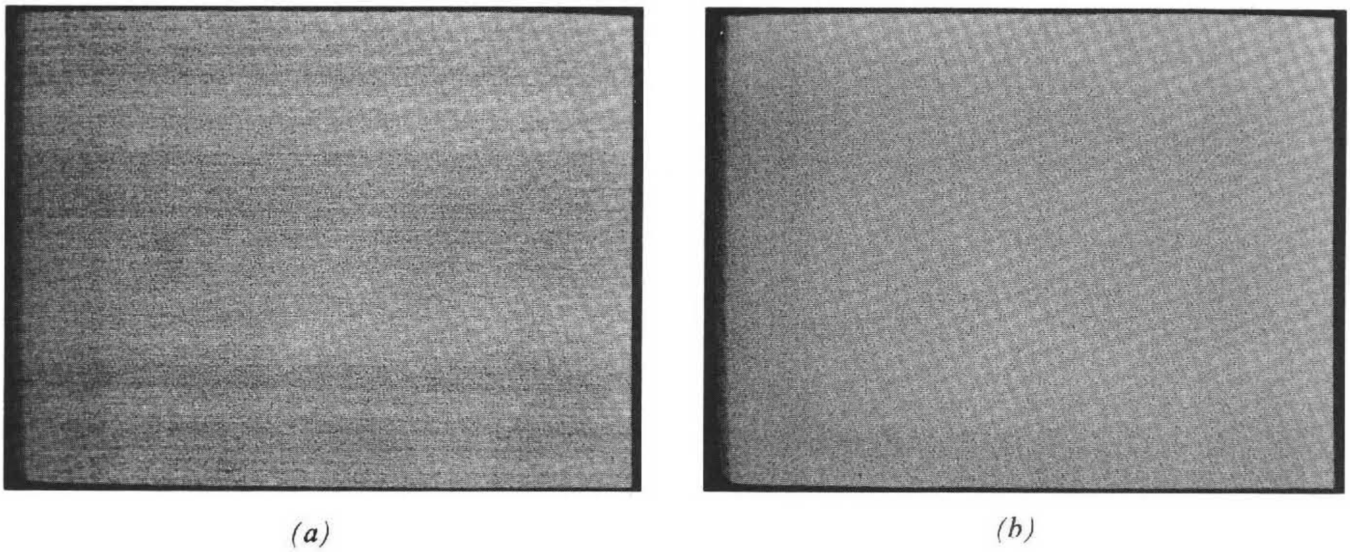
(a)
(b)

Fig. 7 – Effect of noise on television picture, (a) without correlated double sampling, (b) with correlated double sampling

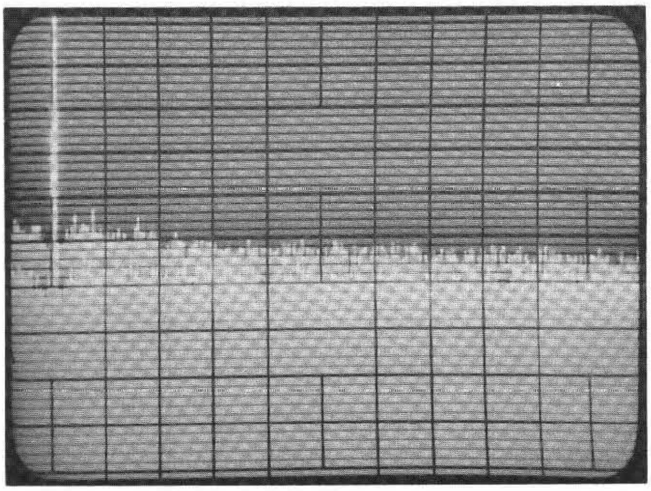

Fig. 8 – Noise spectrum of CCD sensor with and without correlated double sampling. Scales: hor. 10 kHz/div, vert. 10 dB/major div.

feedthrough of the compensation signal. This gave rise to the same type of picture impairment that the circuits were intended to reduce (vertical stripes – but in dark areas of the picture rather than highlight areas). A second drawback was that the stripe stripper was "open loop"; the gain corrections were calculated by reference to the input signal only. Thus any errors or drifts in the concealment system itself remained uncorrected.

An improved stripe stripper was therefore implemented as shown in Fig. 10. An analogue multiplier integrated circuit was used as the core of the system and the incoming video signal was multiplied by a set of stored correction factors; these correction factors were stored in 8-bit digital form and fed to the multiplier via an 8-bit DAC. The gain of the DAC could itself be adjusted – the reason for this will be discussed in section 3.2 of this report. Considerable attention was given to the accuracy of the d.c. levels (set by clamp circuits) of the input signals to the multiplier in order to minimise any feedthrough of the correction signal; in addition the multiplier integrated circuit was temperature stabilised using a small oven (having a warm-up time of approx. $1\frac{1}{2}$ minutes). These precautions were necessary because of the high degree of amplification applied in the subsequent gamma corrector (up to 40 dB for signals near black).

Any remaining slight errors or non-linearities in the stripe stripper were then compensated for by generating the stored correction factors in a closed-loop manner. Under open-gate conditions the output of the multiplier was compared, element by element, with a fixed reference voltage. If the multiplier output for any particular element was less than the reference voltage, then the corresponding correction factor was increased; if the multiplier output was greater than the reference voltage the correction factor was reduced. The correction factors were therefore generated by successive approximation; as well as eliminating the effects of any residual errors in the multiplier, this closed-loop approach was also very efficient. The only circuitry required, other than the stripe stripper itself and its associated control system, was a single analogue comparator. A fast microprocessor, common to all three colour channels, was used for control; this microprocessor also calculated the correction factors, taking approximately $\frac{1}{2}$ second for each colour chanel, and fed the results to the digital storage of the stripe stripper. The corrections were also written into a set of low speed CMOS RAM's whose power supply was maintained with a small, rechargeable battery. Thus the correction factors were not lost when the power was switched off, even for periods of as long as one month.

Long-term accuracy of the stripe stripper system was good and there was no significant drift in

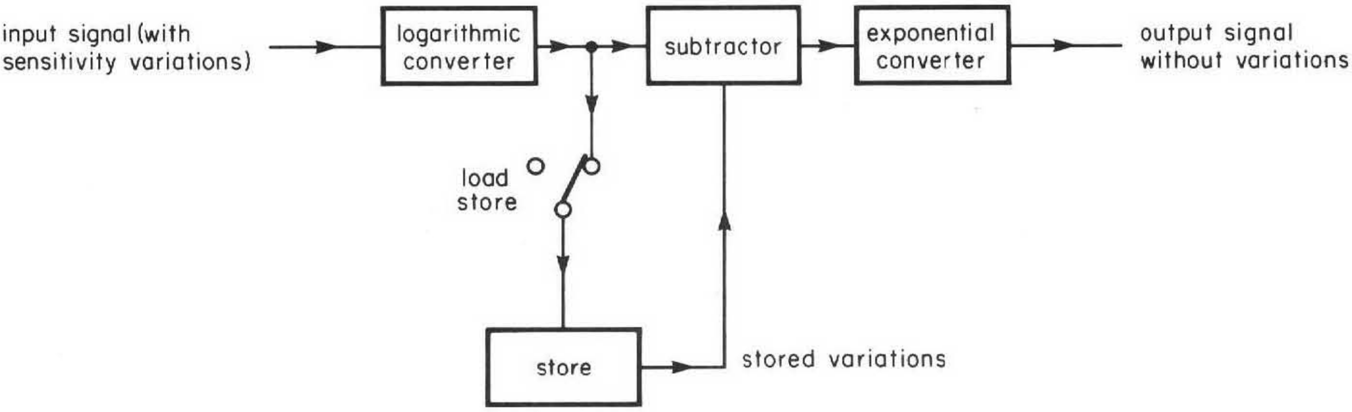

Fig. 9 – Block diagram of original stripe stripper.

the clamp level adjustments. In fact the performance of the telecine was limited, not by the accuracy of the circuits, but by the tendency for the optical system to collect dust; such dust, if present in the film gate, produces an additional coherent pattern of vertical stripes. Thus it was found desirable to orientate the gate so that its tendency to collect dust was minimum; this can be achieved if the film path through the gate is vertical and if any gate lenses are separated slightly from the film itself. A further precaution would be to blow air continuously over any exposed lens surface or to provide a cleaning "blast" of compressed air on demand.

3. Variable speed operation

3.1. Basic control system

There are two characteristics of line-array telecines that make them particularly attractive for running film at non-standard speeds. The first is the predictability of the scanning process and the second is the presence of digital storage at the output. This digital storage is required to convert the signals from the sequentially-scanned form in which they emerge from the sensors into the standard, interlaced format; in addition, however, it may be used as a buffer between the incoming scan rate from the sensor and the outgoing rate locked to station syncs. Thus the sensor scan rate may be allowed to vary in order to track changes in the film speed; for example, if the film were to be replayed at 18 frames per second, the sensor scan rate would have to be reduced to 0.72 times its value for 25 frames per second operation in order to scan the same film area. In addition, if it is necessary to adjust the height of the final television presentation, this can be done by further varying the sensor scan rate. Consequently, the sensor is only rarely scanning at the same rate as the television line standard.

It is possible to achieve a variable scan rate from a line-array sensor in two ways. Either the frequency of the drive signals must be varied, or else the number of clock pulses taken to complete each scan must be altered. Because the output of the sensor is usually lowpass filtered to remove any spurious clock breakthrough, and because it would be very difficult to alter the characteristics of this filter dynamically to track variations in the clock frequency, it is preferable to keep the clock frequency constant. Instead "gaps" of variable numbers of clock pulses are inserted between successive scans; low sensor scan rates are therefore relatively easily accommodated. The maximum scan rate is, however, fixed and so there is a maximum speed at which a film can be run if the sensor is to scan every line. For example, the sensor scan rate required by normal aspect ratio 35 mm film running at 25 frames per second and producing an output on the 625/50 standard is 18.075 kHz (1 line every 55.3 μsec). If a telecine were designed with this as a limit then 16 mm

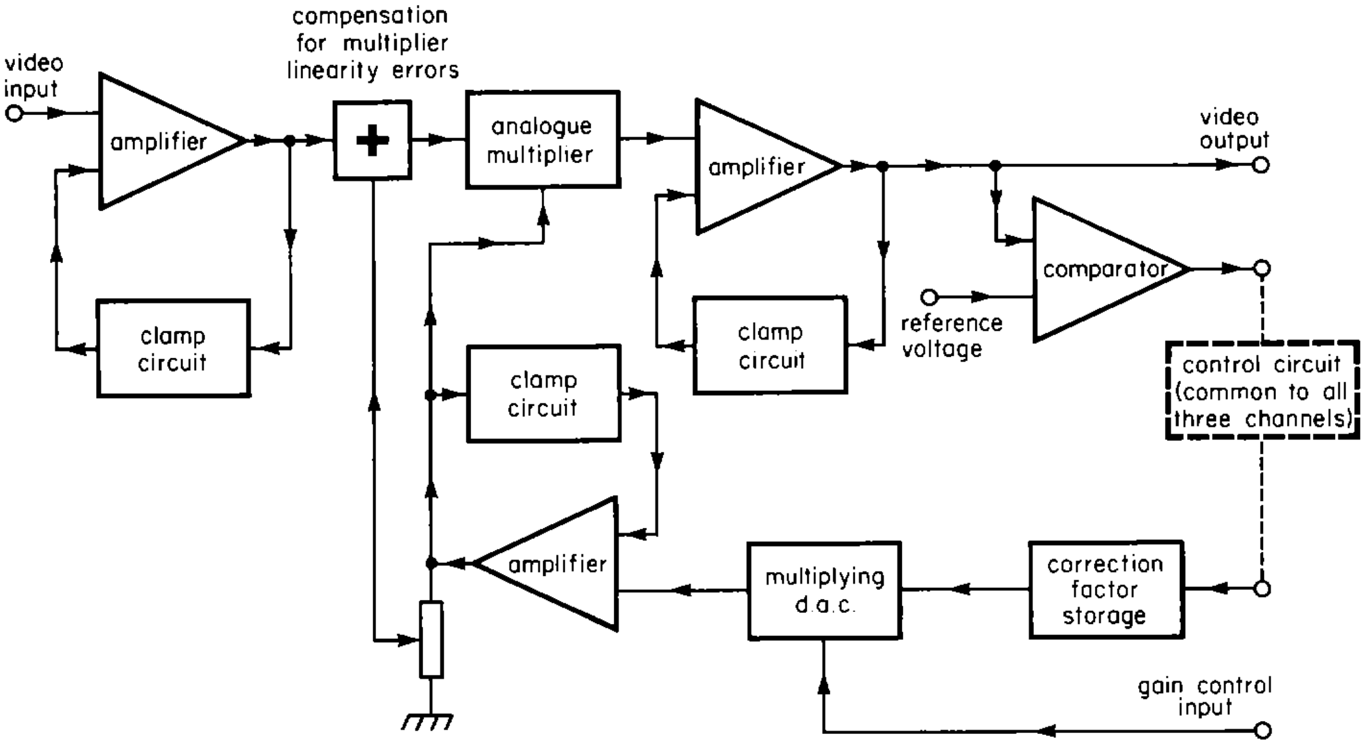

Fig. 10 – Block diagram of improved stripe stripper.

film could be run at speeds up to 28.9 frames per second; any increase in these speeds, or any increase in the displayed picture height, would mean that the sensor could not scan every line. Alternative techniques to overcome this are discussed in section 3.4.

Fig. 11 shows a block diagram of the control system used to link the sensor scanning circuits and the digital store. The heart of the system was a high-speed microprocessor used to control the scan generation circuits for the sensor. This microprocessor measured the speed at which the film was running by timing the interval between two successive pulses produced by a tachometer fitted to the capstan motor driving the film; suitable tachometers could also sense the direction in which the film was being driven. Timing was achieved by counting the number of microprocessor clock cycles, n, between successive tachometer pulses. If the frequency of the microprocessor clock (which was crystal controlled for stability) were f_m Hz then the tachometer pulse frequency would be f_m/n Hz. If there were N tachometer pulses produced for one revolution of the capstan, and if the capstan diameter were d_c metres, then the film would have been moving at a linear speed of $\pi d_c f_m / Nn$ metres per second. And if the distance between successive frames on the film were l_f metres this film speed would correspond to a frame rate of $\pi d_c f_m / N_n l_f$ frames per second. Thus if the height of the displayed part of the film frame were l_a and the output television standard were 625/50, the line scan rate required from the sensors would be f_l, given by the equation:

$$f_l = \frac{575 l_f}{l_a} \frac{\pi d_c f_m}{N l_f} \frac{1}{n} \qquad (3.1.1)$$

A similar equation could be derived for the 525/60 television standard by replacing the factor 575 by the number of active lines in the 525-line standard.

The first factor in equation 3.1.1. (575 l_f/l_a) depended only on the desired height of the final television image and was a user control. The second factor ($\pi d_c f_m N l_f$) was largely constant for any given film gauge; small variations occurred, however, due to variations in the film frame pitch. These variations were compensated by examining at regular intervals the output of an idling sprocket wheel driven by the film perforations. When the sprocket wheel produced a pulse (once per film frame) the particular line being scanned at that instant, L_n, was noted and compared with the corresponding value from the previous sprocket pulse, L_{n-1}. Ideally L_n should be the same as L_{n-1}; a line difference, $L_n - L_{n-1}$ (corrected for any additional factors – for example the number of whole film frames having been scanned between the two sprocket pulses being received) was generated. If this were not zero then clearly the original value assumed for l_f was in error and a new value, l_f', was calculated using the

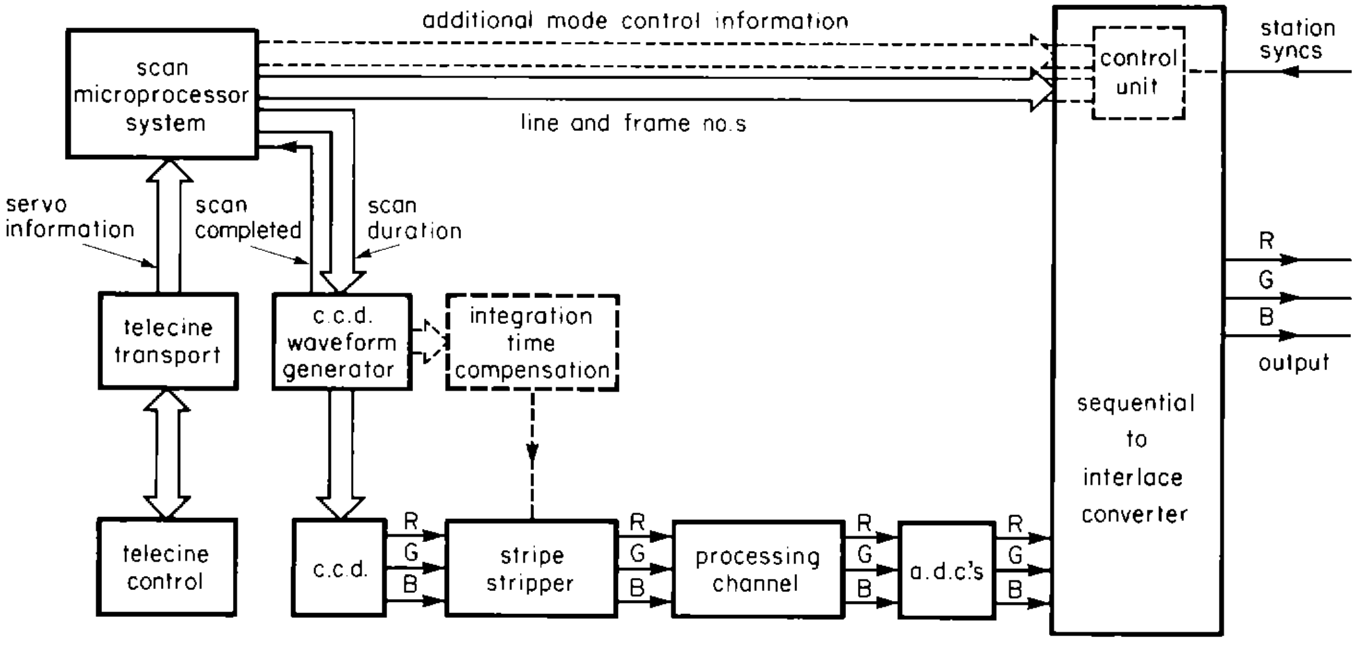

Fig. 11 – Control system for line-array telecine.

formula:

$$l_f' = (1 \pm A(L_n - L_{n-1}))\cdot l_f \qquad (3.1.2)$$

where A was a constant whose value was chosen to achieve a satisfactory compromise between an excessively long response time (A too small) and one excessively sensitive to short-term fluctuations in the sprocket pulse rate (A too large). The plus or minus sign depended on whether the film was running forwards or backwards. A correction was also made for any incorrect film framing by incorporating an additional term:

$$l_f' = (1 \pm \mathrm{A}(L_n - L_{n-1}) \pm B(L_n - F))\cdot l_f \qquad (3.1.3)$$

where F was a number indicating the desired framing relationship of the film.

When the microprocessor had calculated the desired scan rate, information was sent to the waveform generator which produced the signals to drive the sensors. In order to produce these signals, the waveform generator required a stable clock signal at a frequency, f_s, set by the parameters of the sensor itself; such a clock signal was provided by a second crystal oscillator as part of the waveform generator. A convenient way of establishing a given scan rate was therefore to count the number of cycles of this sensor clock oscillator; the information from the microprocessor thus consisted of the desired number of sensor clock cycles in a scan, C_s. This number, C_s, was given by the equation:

$$C_s = \frac{l_a}{575 l_f}\cdot\frac{N f_s l_f}{\pi d_c f_m}\cdot n \qquad (3.1.4)$$

The outputs from the sensors (one for each of the red, green and blue channels) were stripe-stripped, as described in the previous section, processed in the normal way and then written into the sequential-to-interlace store. A "label" consisting of the line and frame number corresponding to every horizontal scan, was sent to the control unit of the sequential-to-interlace store (itself another fast microprocessor); the control unit then used this label to calculate where in the store the particular scan signals should be placed so that they could be retrieved correctly. At the start of every scan the waveform generator sent a flag signal to the scan control microprocessor to enable the line and frame number to be updated; it was also this line number that was used in the calculation of the film shrinkage.

Fig. 12 shows a more detailed block diagram of the sequential-to-interlace converter. The main digital store was fed with a sequence of read and write addresses and incoming data. The incoming data was written, sample by sample, into the storage location specified by the write address. At the same time data was also read from the locations specified by the addresses from the read address generators (only two of which are shown in Fig. 12 – a third was provided for use in connection with the dirt concealment to be described in section 4). The address generators were sent the address of the storage location of the first picture element in any line by the control microprocessor; they then incremented this address, element-by-element, for the duration of one television line. After this time they stopped and a further start address was required from the control microprocessor in order to enable operation to continue. If a start address was not received, or if the control microprocessor generated an invalid start address, the address generator produced a "disable" signal used to inhibit writing to the store (write address generator) or blank the relevant data output (read address generator); the consequent control of signal blanking was useful and an application is described in section 3.3.

The two outputs from the store were used to feed an interpolator which combined the signals in a proportion determined by a control signal generated by the microprocessor. It could be used to provide movement interpolation or to interpolate between lines for high speed running (see section 3.4). The output from the interpolator was then converted back into analogue form.

The fast operating speed of the microprocessor used to control the sensor scanning and the sequential-to-interlace converter allowed a very fast response of the overall control systems to small fluctuations in film speed. Thus, once the correct film shrinkage had been acquired, the scanning system ensured that broadcast-quality signals could be produced from film running at any speed up to the limit set by the sensor clock frequency. In the future this should allow telecine machines to be started "on-air" without any objectionable disturbances on the output.

3.2. Compensation for integration time variations

With the control system described above, the sensor scan rate depended on the running speed of the film; in turn this caused the integration time of the sensors to vary producing an apparent variation in their sensitivity. If this had remained uncorrected, it would have caused unwanted fluctuations in the gain of the final output. Fortunately, it was relatively simple to supply an additional gain-controlling signal to the analogue stripe-stripper; this signal adjusted the gain of the DAC carrying the stored correction factors in order to compensate for

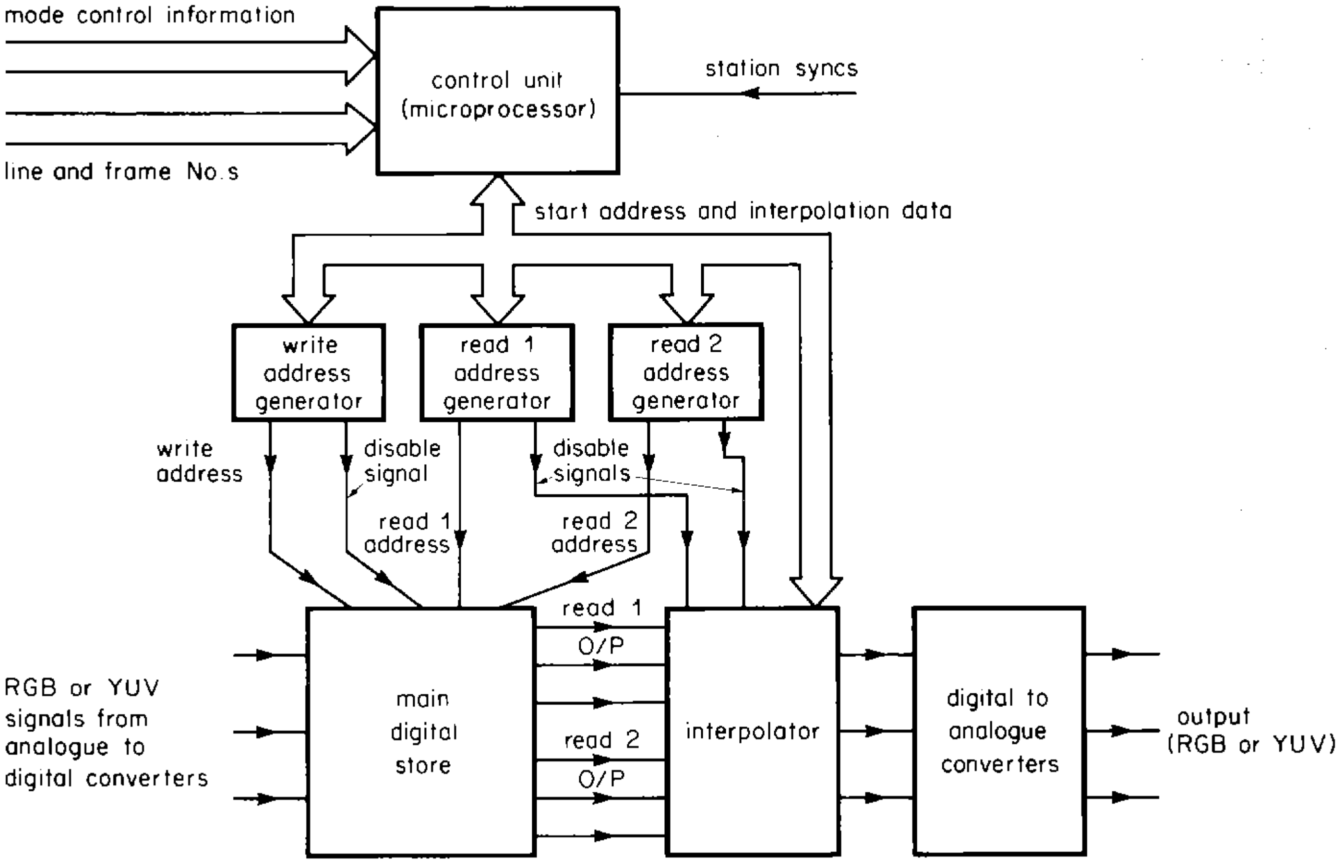

Fig. 12 – Block diagram of sequential-to-interlace converter.

line-to-line variations in integration time. The signal was generated from the scan rate information fed to the waveform generator (as shown dotted in Fig. 11); as an alternative it could relatively easily have been calculated directly by the scan microprocessor.

3.3. Operation at slow film speeds

Although it would have in theory been possible to use the integration time compensation system of the previous paragraph for the complete range of film speeds down to very nearly still frame, there were practical problems with doing so. Chief among these was the danger that the sensor could overload at very long integration times. In order to combat this problem the scanning frequency range of the sensor was restricted to 2:1. If the scan rate fell below one half of the maximum possible, two scans were generated and one of these discarded; such dummy scans were easily generated by a suitable program in the scan microprocessor and were identified as non-valid by sending an additional signal to the waveform generator (an alternative would have been simply to have omitted the label sent to the sequential-to-interlace converter).

As the film slowed down even further, more and more dummy scans were generated so as to keep the average integration time of the sensors within the 2:1 range. In order to minimise the visibility of any errors in the compensation process, the duration of all the valid scans was kept constant and the number and duration of the dummy scans varied in order to track film speed changes.

The output of the telecine, at slow film speeds, consisted of a number of fields from one film frame followed by a number from the next film frame, and so on. This correctly gave the impression that the movement in the original scene was slowing down. However, the film could still be running at several frames per second and it was surprisingly difficult to locate a specific event (a shot change, for example) with confidence. Since one of the main applications for slow speed running was considered to be for previewing film and setting up pre-programmed colour corrections, an alternative method of presentation was developed. The display was made to "wipe" vertically from one frame to the next; the join between the two scenes showed the current line being scanned in the gate and was marked with a black line to enable the junction to be identifiable even when both frames were similar. By this means it was possible to locate a given event very confidently (see Fig. 13) – it has even proved possible to "rock" the capstan gently backwards and forwards by hand in a similar way to that used more conventionally in audio recording.

3.4. Operation at fast film speeds

As has been mentioned above, there was a basic difficulty when film was run at significantly above normal speeds – this was the fact that the sensor could no longer scan fast enough. Thus not every line could be scanned on every film frame. Two solutions were tried. The first assembled the television picture over several film frames; for example, if the film were moving at four times normal speed (approximately 100 frames per second), the first television field would be composed from lines 1, 5, 9, 13 etc. of film frame one interleaved with lines 3, 7, 11, 15 etc. of film frame two. The following (interlaced) television field would contain lines 2, 6, 10, 14 etc. of film frame three and lines 4, 8, 12, 16 etc. of film frame four. The correct temporal sequence of the film frames was thus maintained, and the vertical resolution of the television picture was marred only by the increased vertical movement of the film during the integration time of the sensors; however, the picture split up into a multiple image as soon as there was any movement in the scene. This made the method undesirable for most of the time.

The second method was found to produce far more acceptable results under most conditions. If the film was again moving at four times normal speed, the first TV field would be composed from lines 1, 5, 9, 13 etc. of film frame one only; the missing lines would be generated either by simple repetition of the scanned lines or by interpolation between them This repetition or interpolation produced a further loss of vertical resolution and, because not every film frame was scanned, it was more difficult to locate events lasting for only a few film frames. Nevertheless, the results were quite adequate for review purposes; Figs. 14(a) and 14(b) show pictures produced at a film speed of 10 times normal by the duplication or interpolation of every sixteenth line from the scanned image.

There are cases, however, where broadcast quality signals (in contrast to signals for review purposes) might be required. For example, it might be desirable to replay film at speeds just higher than normal (up to 30 frames per second, for example) or to increase the displayed height of film running at 25 frames per second (using 1.85 : 1 aspect ratio film, perhaps). Under these conditions, where only every other line is available, it is clearly important to improve the quality as much as possible. The loss of vertical resolution, under these conditions, was not very severe and could have been recovered by suitable aperture correction; this aperture correction would need different characteristics from more conventional correctors, the only drawback being an increase in the high frequency components of the sensor noise (*not* the film grain). To some extent this

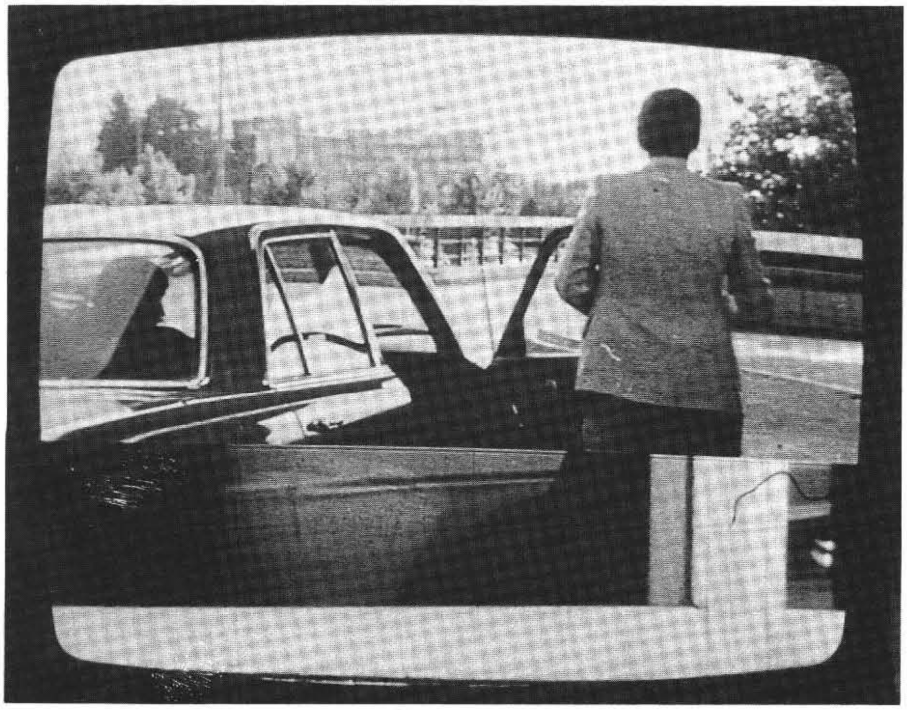

Fig. 13 – "Wipe" presentation of slowly moving film demonstrating the easy visibility of shot changes.

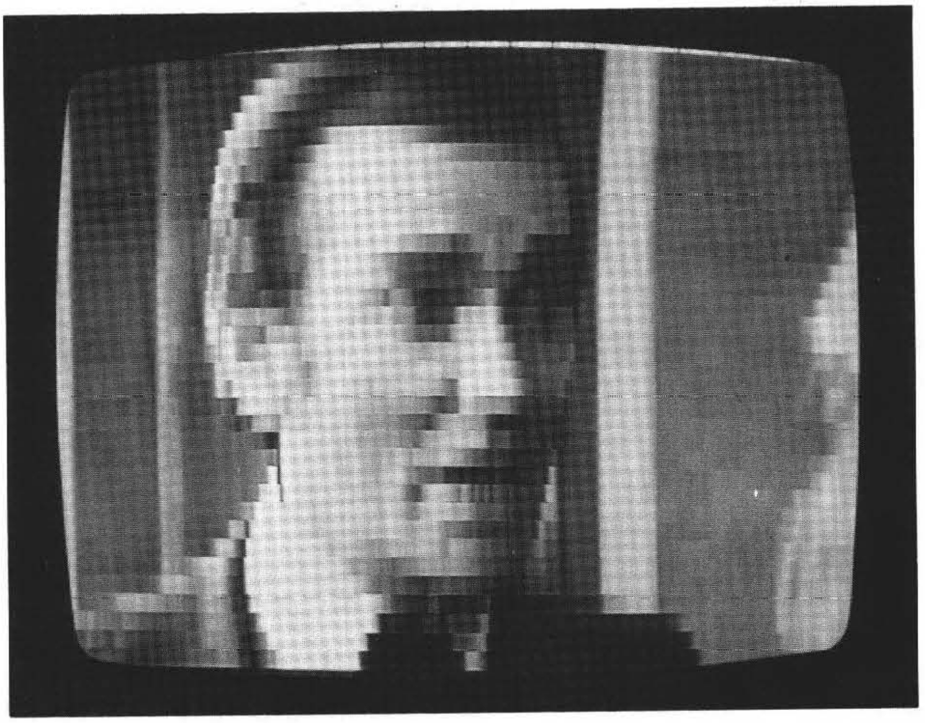

Fig. 14 – Television output from film running at 250 frames per second.

(a) unscanned lines duplicated

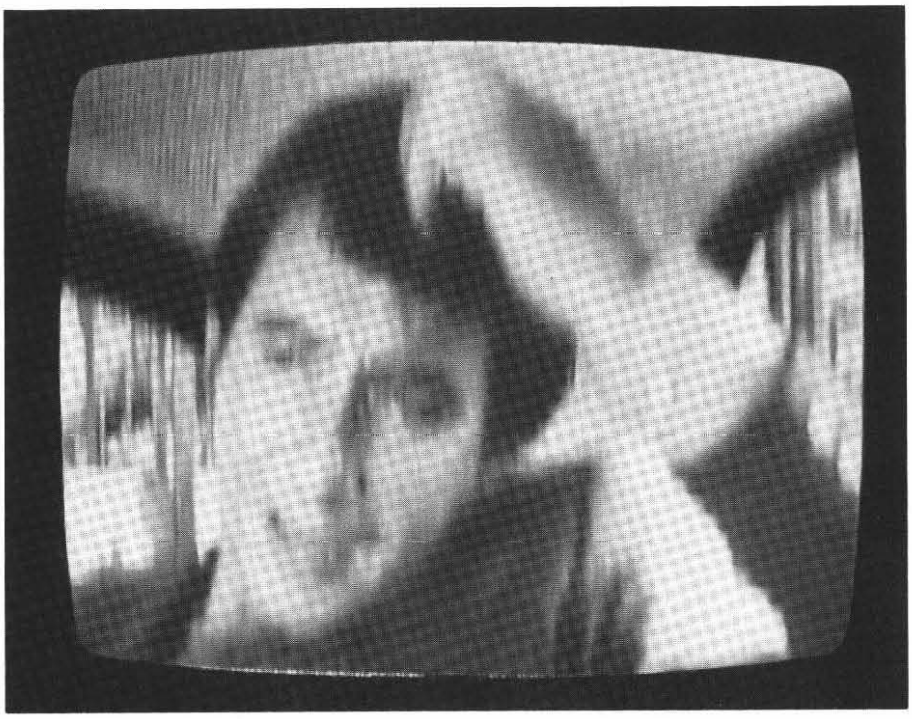

(b) unscanned lines interpolated

increase in noise would be offset by the increase in signal gained by the integration of light over two lines on the film. A more difficult problem was to make the alias components, produced by scanning only 312 lines, have as low a visibility as possible. This was achieved by a combination of the optimum parameters for interpolating the additional lines and by not always scanning the same lines; for example, on the first film frame lines 1, 3, 5, 7 etc. were scanned and lines 2, 4, 6, 8 etc. interpolated, and on the second film frame the situation was reversed. The aliasing was thus broken up and concealed by the normal aliasing produced by the television scanning process. Broadcast quality pictures were produced at speeds up to more than 50 frames per second (Fig. 15). Indeed, the results from film running at up to

100 frames per second (Fig. 16) were still surprisingly good, although aliasing resulting from the scanning of only 156 lines was clearly visible.

Implementation of this fast-running mode in the control system of Fig. 11 was relatively simple. The scan microprocessor examined the desired number of clock pulses per line scan, C_s. If C_s was smaller than the number of clock cycles required to complete the readout of all elements of the sensors, then a decision was made to scan only a proportion of the total number of lines. C_s was then doubled and re-examined; this sequence of doubling and examination was repeated until it was possible to achieve a complete scan in the required number of clock pulses. Thus, in order to keep the amount of ad-

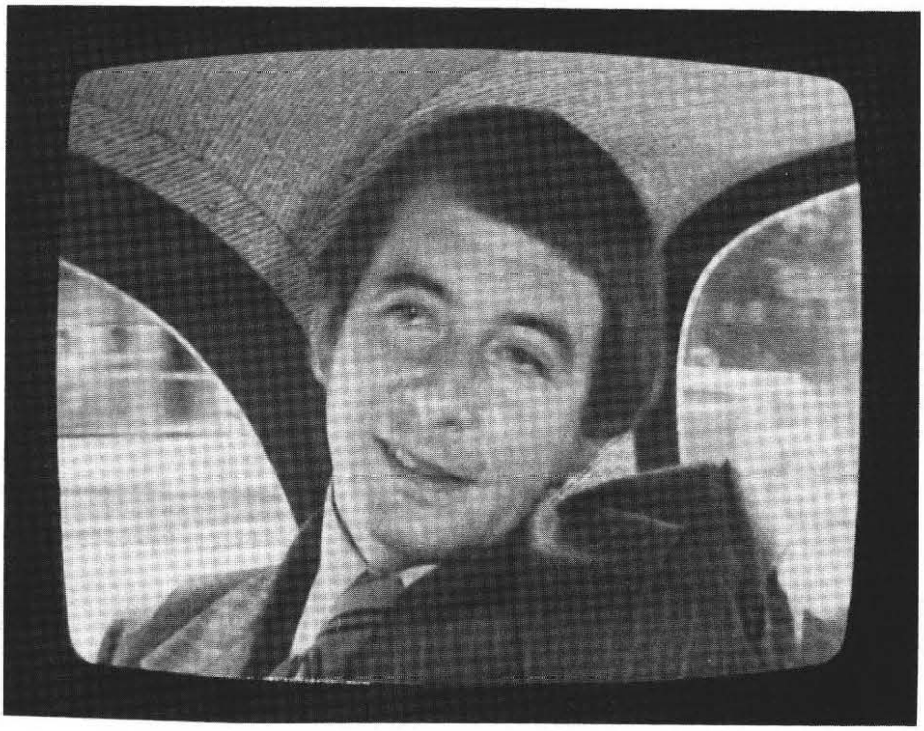

Fig. 15 – Television output from film running at 50 frames per second.

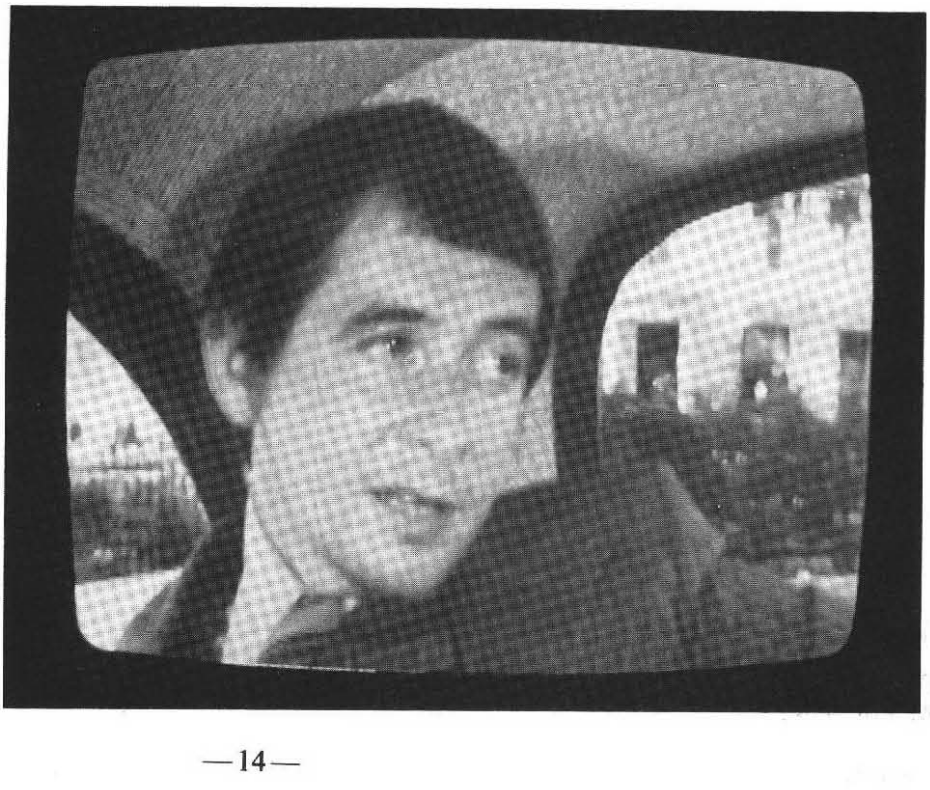

Fig. 16 – Television output from film running at 100 frames per second.

ditional processing as small as possible, the proportion of the lines scanned was reduced in a binary sequence (i.e. either a half, or a quarter, or an eighth, etc., depending on the speed at which the film was running). In order to enable the sequential-to-interlace converter to reconstruct the pictures correctly, information about this proportion was sent to the sequential-to-interlace control microprocessor; this connection is shown dotted in Fig. 11.

3.5. Movement interpolation

One of the benefits of the system described in this Report is the ability to run film at any speed. However, this can result in problems with movement portrayal. The system of Fig. 11 generates n interlaced television fields from m sequentially-scanned film frames, where m and n are unrelated numbers. This is achieved by repeating each film picture n/m times. Where n/m is not a whole number the repetition rate may cycle. For example, $n/m = 2\frac{1}{2}$ can be achieved by repeating the first frame twice, the second three times, the third twice and so on. Unfortunately, this leads to unacceptable judder on moving scenes; to combat this movement interpolation can be used.

Fig. 17(a) illustrates the situation. The film pictures are samples of the initial scene at discrete moments in time, shown by the circles; each circle shows the time at which a given line was sampled. It is necessary to generate an approximation of what these samples would have been if the scene had been recorded by a television camera. The crosses in Fig. 17(a) correspond to the time at which each line would have been scanned in this case. The slope of the television fields in time is a complication which can be neglected with very little penalty; Fig. 17(b) shows the situation if this is done. A simple movement interpolator would generate an approximation to the signal that would have been generated by a television camera by adding together a proportion of film frame 1 to a different proportion of film frame 2. If these proportions are S_1 and S_2 respectively then, in order to avoid flicker in the final picture:

$$S_2 = 1 - S_1 \qquad (3.5.1)$$

Several laws that were tried relating S_1 (and hence S_2) to τ (the relative time relationship between television fields and film frames – see Fig. 17(b)) are shown in Fig. 18. Law A corresponded to the case where there was no movement interpolation. In contrast, Law B had an excessive amount of movement interpolation; no judder was visible on movement when this law was used, but instead there was a pronounced loss of resolution in moving areas. These were an infinite number of compromise laws (such as Law C in Fig. 18) which returned some of the lost resolution at the expense of re-introducing some movement judder; in general, for most movement speeds the optimum interpolation characteristic was judged to be close to Law C as shown, and this was judged to be a significant improvement over Law A or B.

There was no need to restrict the law to a straight line (Law D is shown to emphasise this point) although in practice little was gained by a departure from such a characteristic. Likewise there was no need for the interpolation characteristic to remain unaltered for all film speeds; the control system of Fig. 11 would allow the interpolation parameters to vary according to different film speeds or television standards so as always to operate the telecine under optimum conditions.

4. Dirt concealment

One of the impairments visible on programs originating on film is the occurrence of blemishes on the picture due to the presence of dirt. The blemishes may be produced at any one of several stages.

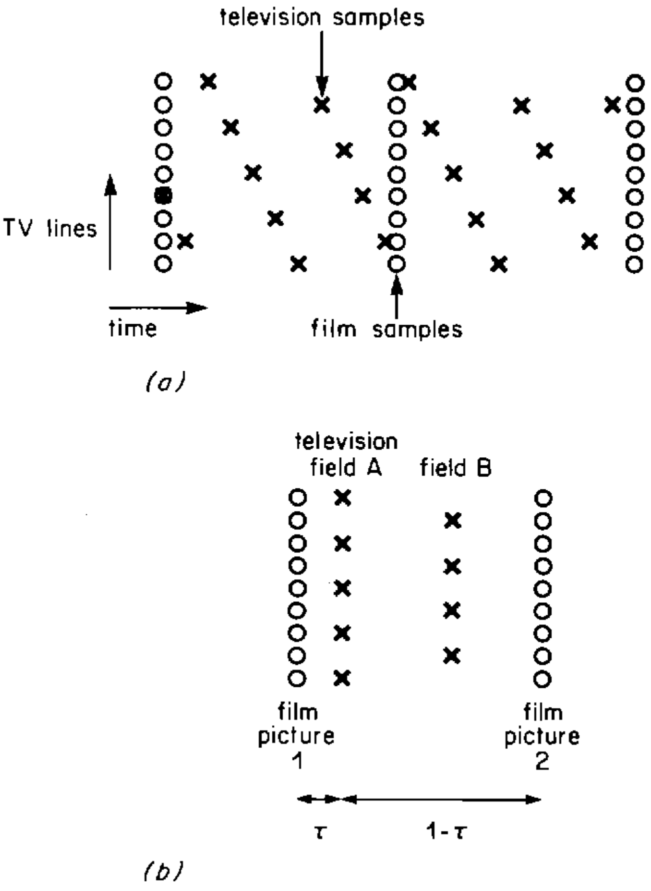

Fig. 17 – Temporal relationship of film frames and television fields.

1. Dirt may be deposited on the film as it runs through the camera. This dirt then prevents the light from exposing the film; the area of the film that was covered by the dirt then appears black on the final televised image.

2. If negative film stock was used in the camera and it is not considered desirable to run the original negative in the telecine, a print must be made. Any dirt deposited on either negative or print stock before printing prevents the print film being exposed. This then produces a blemish that appears white on the final picture.

3. Dirt may be deposited on the film before it is run in the telecine. The appearance of this dirt on the final picture then depends on the type of film. Dirt on negative film appears white, while dirt on print or reversal film appears black.

An earlier investigation into the characteristics of dirt on televised film yielded three important conclusions.[12] The first of these was that the amount of dirt on a piece of film increases considerably with handling.

A second conclusion was that dirt occurring at negative stages (that is, dirt appearing white on the final picture) is very much more objectionable than dirt occurring at positive stages. A difference of 10–20 times in the number of dirt particles giving rise to a given subjective impairment was noted. Balancing this finding, however, is the fact that a film negative is usually handled more carefully than a print.

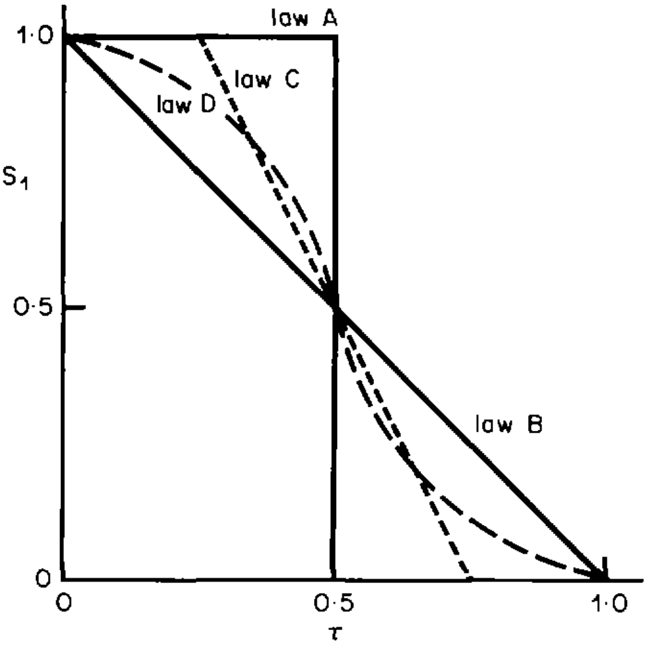

Fig. 18 – Possible interpolation laws.

The final important conclusion was that dirt particles are usually very small; more than half of the particles are less than 25μ in diameter (this size on 16 mm film corresponds to a blemish lasting for only 130 nsec or less than 2 sample periods of the CCIR digital standard). Only 1% of all dirt particles were larger than 55μ (280 nsec for 16 mm film, or just under four samples of the CCIR standard). This has a bearing on the choice of a suitable concealment algorithm.

4.1. Detection of dirt by infra-red light

It had been known for some time that dirt physically present on the final film (as distinct from the image of dirt present at an earlier stage) may be detected by examining the film with infra-red light.[13] This is because colour film dyes are largely transparent to infra-red light. The practical application of this knowledge, had however, been restrained by the difficulty of producing such an infra-red signal without incurring other problems. For example, flying-spot tubes with short persistence afterglow characteristics do not emit any significant amount of infra-red light. It is possible to incorporate a mixture of normal and infra-red emitting phosphors, but this causes severe problems of increased phosphor grain in the visible channels. Line-array telecines, on the other hand, are ideally suited to this form of dirt detection. Not only do the sensors have a response extending to wavelengths beyond 1 μm, but the light source used (a quartz halogen lamp) is rich in infra-red output.

Figure 19 shows the transmission characteristics of two samples of fully exposed colour print film as a function of wavelength. It can be seen that the film dyes become much less opaque once the wavelength becomes longer than 800–850 nm. If a fourth, infra-red sensitive channel is added to the colour splitter assembly, its output is much less

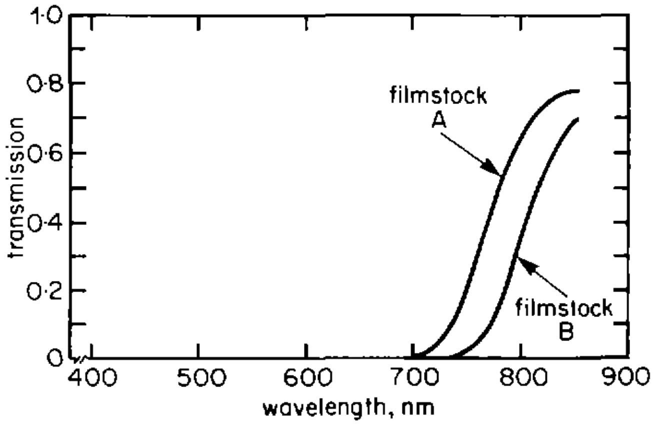

Fig. 19 – Transmission of colour film to infra-red light.

dependent on the scene content and much more dependent on the amount of dirt on the film. Figure 20 shows the optical arrangement used for the experiments. Infra-red light was first separated from the visible components by a dichroic mirror, the visible light then passed through a colour splitter block onto three sensors in the conventional way, while the infra-red light passed through a further shaping filter and was incident on a fourth sensor. Figure 21 shows a photograph of this experimental four-channel analysis system.

The output from the infra-red sensor is shown in Fig. 22, after conversion from sequential-to-interlaced form. The image of the original scene is very much reduced in visibility and the dirt on the film is clearly visible. Further reduction of the picture image contrast is possible by simple masking derived from the visible red channel, since only the cyan film dye has a significant effect on the infra-red transmission. After masking, the signal is passed to a threshold detection circuit (as in the block diagram of Fig. 23) whose output indicates the presence of dirt on the film.

Two further points deserve mention. The first is that dirt cannot be detected until the infra-red transmission has fallen to less than some fixed proportion of its peak value. Thus the extreme edges of dirt particles are not detected. The detected dirt must therefore be expanded by some fixed amount[14] to prevent the appearance of dirt "halos" in the concealed picture.

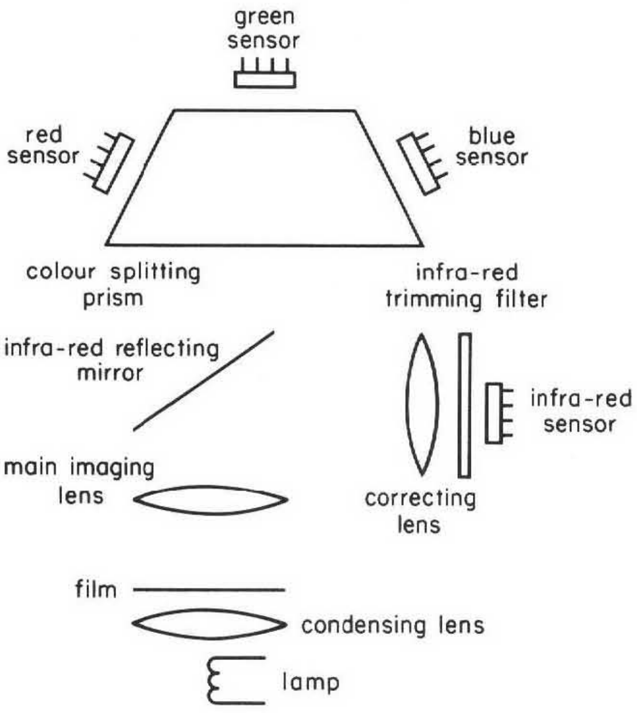

Fig. 20 – Optical arrangement of experimental four-channel analysis system.

The second point is the unexpected enhancement in the infra-red channel of edge detail in the original scene; this edge enhancement is clearly visible in Fig. 22, and is unaffected by masking. It is caused by the poor focus performance in the infra-red of the main imaging lens used in this experiment and this effect is shown diagramatically in Fig. 24. Although the film dyes are transparent to infra-red light, they still display a thickness variation depending on the visible image; this thickness variation is visible as a relief image on the emulsion side of film. The edge of each piece of detail in the scene thus appears, in cross-section, similar to a tiny prism.

If the film emulsion is exactly focused on the infra-red sensor, then the presence of this relief image does not matter. If, however, the sensor is focused on some other plane (i.e. the plane a–a′ in Fig. 24) some of the light from one side of an edge is refracted to the other side, giving rise to the appearance of a dark and a light fringe around each object. The magnitude of the dark fringe affects the closeness to which the threshold for dirt detection can approach the signal level for clear film, and consequently the performance of the dirt detection system on very small particles of dirt. It is thus important to keep the infra-red sensor image in focus, and to use a lens that is free from aberrations at infra-red wavelengths.

The lens used in the experimental system was not designed for use in the infra-red and showed

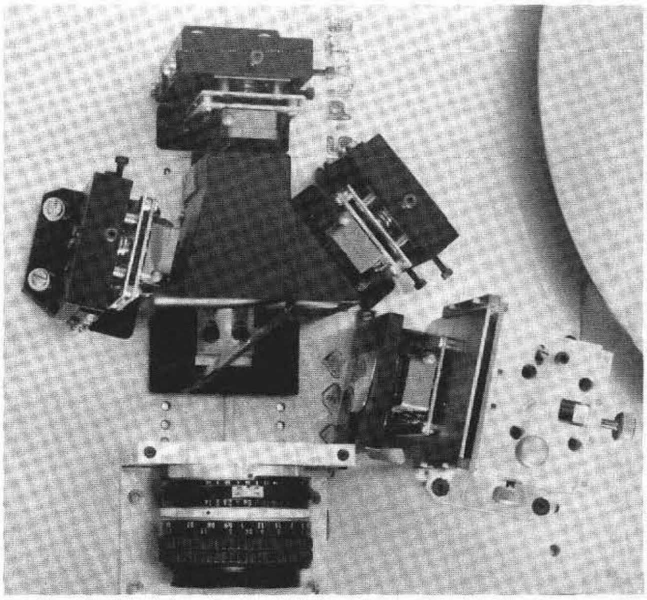

Fig. 21 – Photograph of experimental optical arrangement.

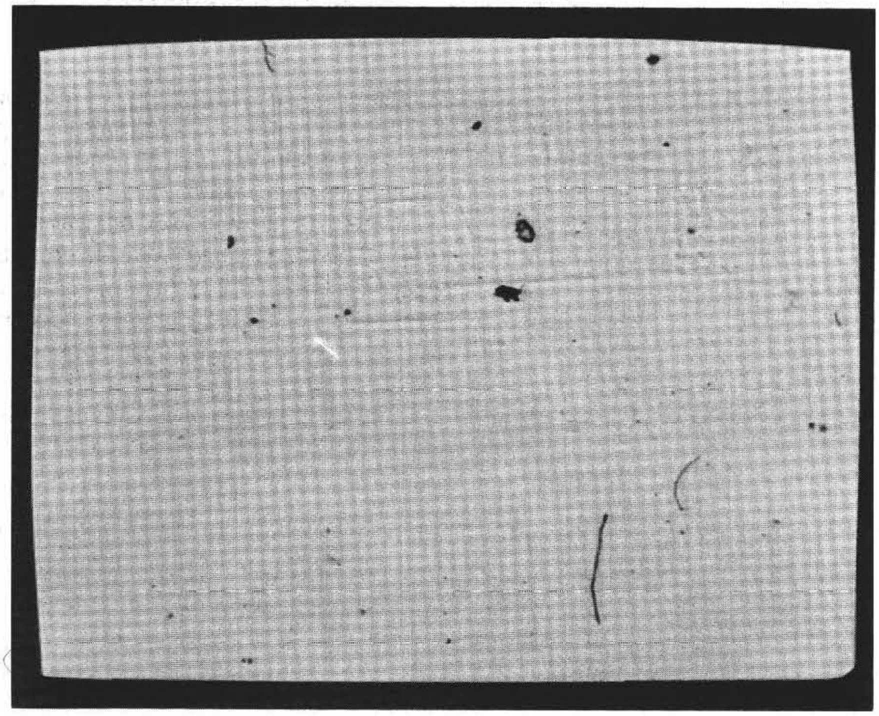

Fig. 22 – Output of infra-red channel.

considerable aberrations. These were minimised by the use of a secondary correcting lens (also shown in Fig. 20), but they were still a limiting factor on the performance of the detection system.

In spite of these limitations, the experimental system was still very successful at detecting physical dirt. Scratches could also be detected as the light falling on the edges of a scratch appears dark in infra-red light.

Another limitation of the detection system is that it fails on monochrome film because the silver image is opaque to infra-red light. The failure to detect dirt under these conditions is unfortunate, and it is important to inhibit the spurious detection of the image itself as dirt. Consequently, circuitry was included in the experimental system to inhibit the dirt detection circuits in the presence of monochrome film.

4.2. Concealment strategies

Once the positions of the blemishes on the film have been detected, they must be concealed by replacing the signal in the region of the blemish with a signal derived by other means. In almost all cases this means estimating what the picture information concealed by the blemish actually is, based upon measurements taken on information in a different part of the scene. Thus there is considerable room for error, and concealment strategies should be chosen not only for the best performance in establishing a good concealing signal on the highest proportion of blemishes, but also for the least obtrusive errors on picture material that causes the strategy to fail. Several of the possible strategies will now be considered in more detail.

4.2.1. Variable-gain approach

As has already been noted, most dirt is very small. Also the dirt deposited on the side of the film away from the emulsion will be out of focus. In neither of these cases will the signals in the visible channels fall completely to black being limited, in one case, by the bandwidth in the imaging system and, in the other case, by the depth of field of the main lens. If the infra-red channel could be relied on to give an accurate measure of the optical attenuation applied to the visible signals, a compensating amount of electrical gain could be applied to conceal the dirt.

The advantage of this approach* is that it potentially gives an output signal that is exactly the same as it would be if the film were not dirty. There are several pitfalls, however. One is that it relies heavily on the resolution and registration performance of the infra-red channel. Bearing in mind that the performance of both the lens and the CCD sensor itself falls off at longer wavelengths, it seems unlikely that perfect compensation could be achieved. In addition, the system would work best at eliminating the dirt that is least visible. In the

*Suggested by Dr J. O. Drewery.

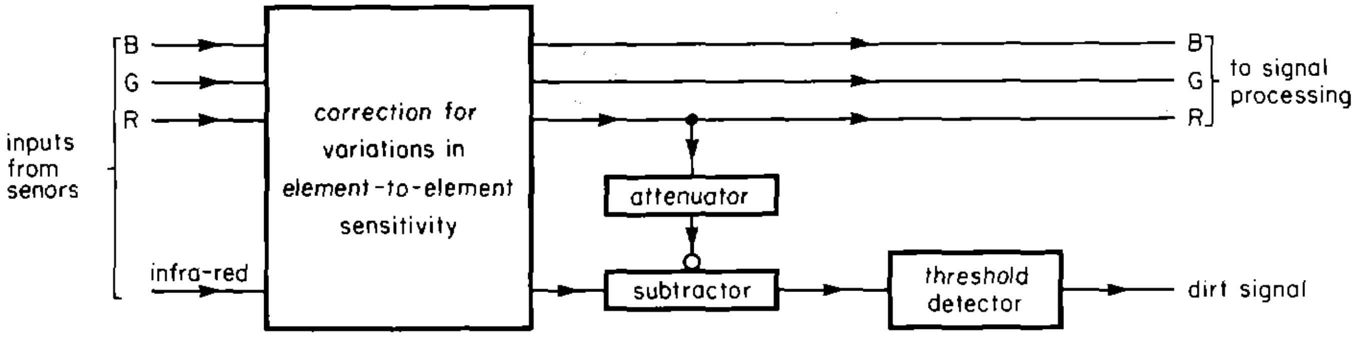

Fig. 23 – Infra-red channel signal processing.

presence of large particles of dirt, the system would fail and an alternative strategy would need to be used. For these reasons, this approach was not pursued.

4.2.2 Concealment by substitution

All other concealment strategies rely on replacing the corrupted picture information by information taken from a different area of the same piece of film. For example, picture elements or lines adjacent to the blemished area may be used or the corresponding area of an adjacent film frame. If a substitution is effected by using adjacent areas of the same film frame, errors can occur in regions of high detail; if a previous film frame is used, errors can occur in moving areas. Most strategies can also be implemented either in recursive or non-recursive fashion, the difference being that, in a recursive implementation the concealed signal itself is used as a basis for generating further concealments whereas, in a non-recursive implementation the concealing signal is generated from the raw incoming picture information; Figs 25(a) and (b) shows the block diagrams of a recursive and a non-recursive dirt concealer based on the use of previous frame substitution to illustrate the difference.

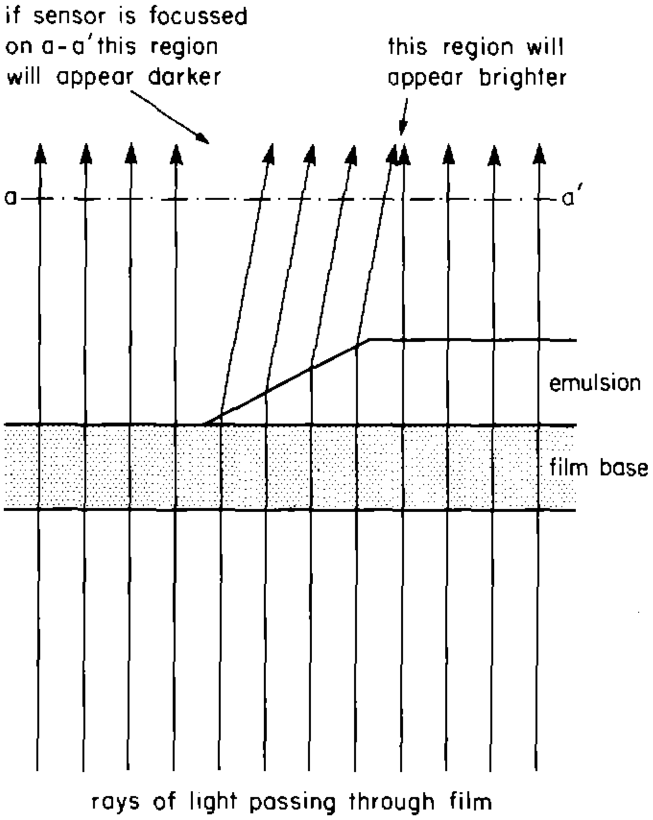
if sensor is focussed on a-a' this region will appear darker
this region will appear brighter
a
a'
emulsion
film base
rays of light passing through film

Fig. 24 – Mechanism for edge enhancement in the infra-red image.

In general, recursive methods are more efficient in their use of storage, but more subject to errors. For example, in the previous-frame concealers of Figs 25(a) and (b), if dirt repeatedly occurs in the same place on many film frames (this could be caused by a vertical scratch, for instance) the recursive concealer of Fig. 25(a) would continue to replace the corrupted material with material taken from the last unblemished frame; the non-recursive concealer of Fig. 25(b) will, instead, substitute the corrupted information from the frame immediately prior to the current output frame. While Fig. 25(a) will give the better performance on stationary scenes, errors caused by subsequent shot changes or scene movement could persist for a very long time. For this reason, it was decided at the outset to investigate non-recursive methods of concealment and to use recursive methods only if significant advantages could be predicted.

4.3. Experimental equipment

Fig. 26 shows a block diagram of the equipment assembled for tests of different non-recursive concealment algorithms with the concealment being effected at the output of the sequential-to-interlace store. In order to avoid the need for an additional store to contain the detected dirt signal, the dirt signal was combined with the video information. Any picture element detected as being corrupted by

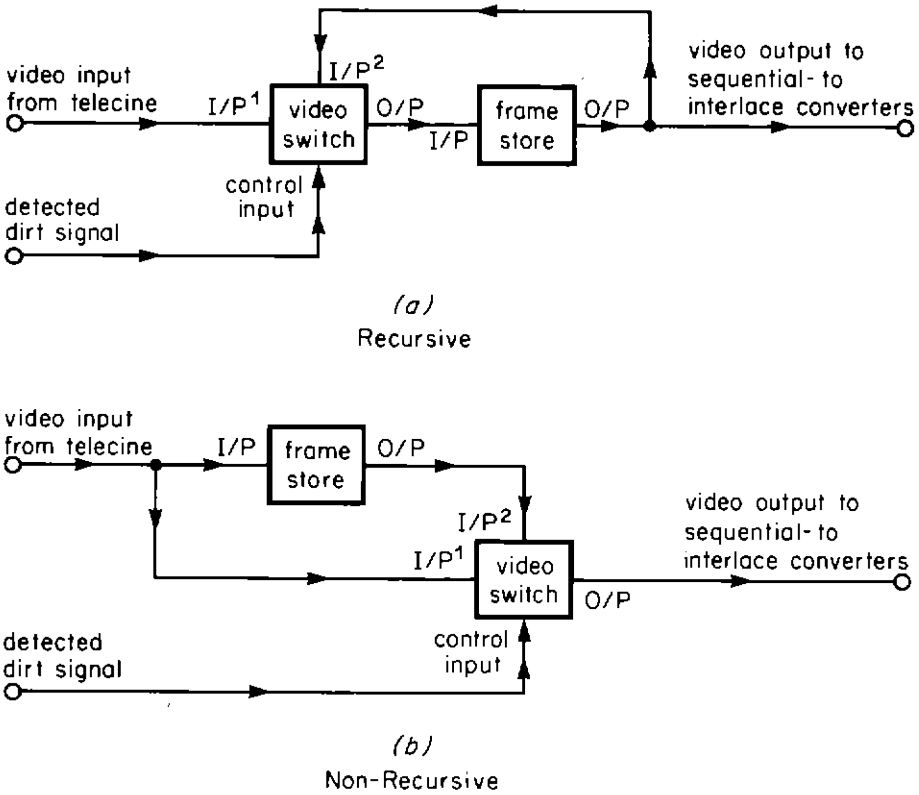

Fig. 25 – Dirt concealment using previous-frame substitution.

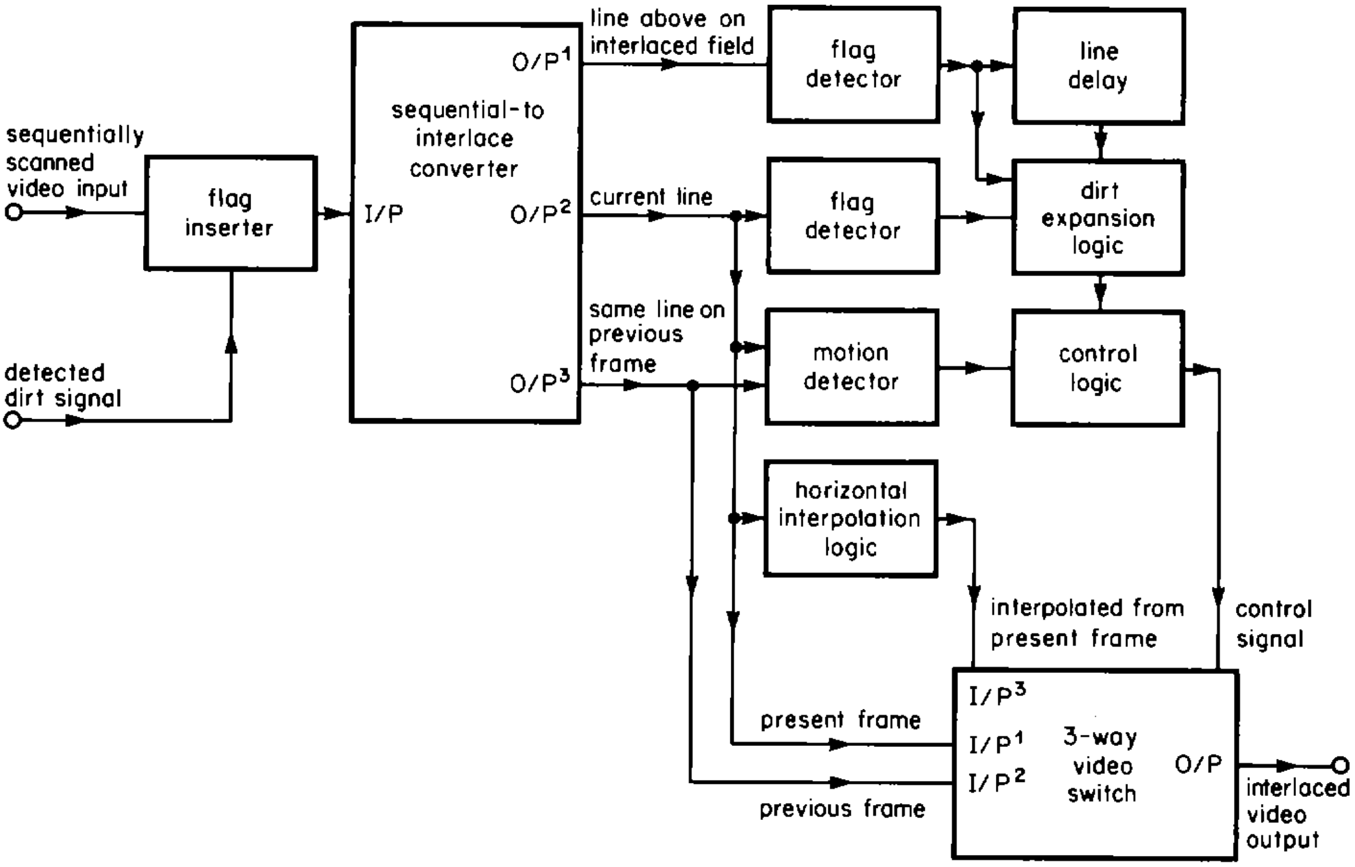

Fig. 26 – Experimental dirt concealment system.

dirt was flagged as dirty by forcing the digital sample to the binary code zero. This was well below the normal telecine black level (set to binary code sixteen) and so enabled the blemished samples to be easily detected at the store output.

The concealing circuitry was deliberately kept very simple to allow several different techniques to be explored without undue waste of effort. Because of this, operation was restricted to monochrome only and the maximum size of dirt that could be concealed was set as sixteen successive samples. In addition, if dirt particles occurred separated by less than sixteen samples, only the first would be concealed. Diagnostic facilities were, however, included to identify those particles of dirt that could not be concealed because of the restricted design of the equipment; this was in order to separate the instrumental limitations from those of the basic algorithm being studied.

The first algorithm to be investigated used replacement of the blemished areas of the picture by the corresponding areas of the previous film frame. This produced a very satisfactory concealment (see Fig. 27(a) and (b) which show, respectively, a picture before and after concealment) for stationary scenes. Scenes containing movement were, however, found to give rise to very visible errors (see Fig. 28), although the majority of the dirt particles were still being satisfactorily concealed. Nevertheless, the objectionable nature of the artefacts on the few blemishes being erroneously concealed was such as to render this particular concealment algorithm useless without some form of protection against scene movement.

A motion detector was therefore incorporated which examined the difference signal between the frame being concealed and the previous frame; in order to prevent the dirt itself from being detected as

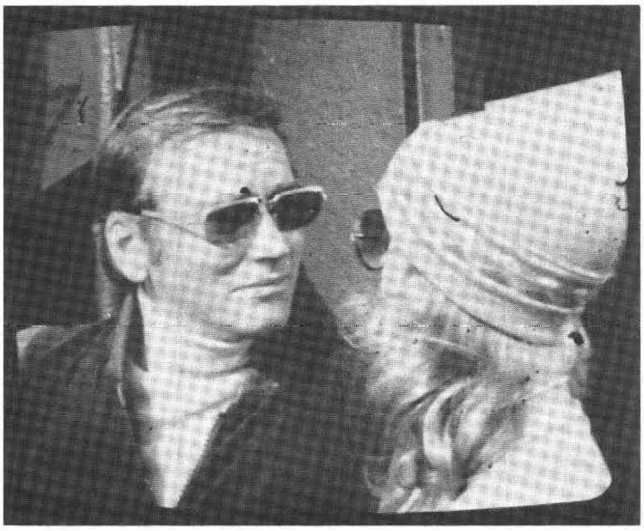

(a) dirt concealment inhibited, showing areas corrupted by dirt

Fig. 27 – Performance of previous-frame substitution on a stationary scene.

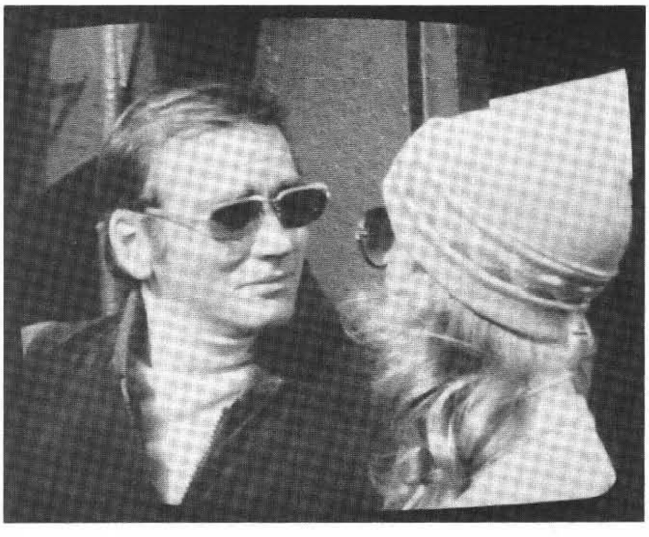

(b) dirt concealment enabled

motion, only information immediately adjacent to the blemished picture elements was used (see Fig. 29). If this difference signal exceeded a given threshold, dirt concealment was suppressed. The performance of the motion detector was not good, owing to some of the simplifications made in its design, and noise in the video signal due to film grain often caused stationary parts of the scene to be wrongly detected as moving. Nevertheless, some improvement in the overall quality of concealment was noted. Fig. 30(a) illustrates this; it shows the same film frame as Fig. 28(b) but with protection against erroneous concealment due to film motion. The major artefacts near the edge of the moving apron are removed and the original dirt allowed to reappear. It is difficult to gain an adequate

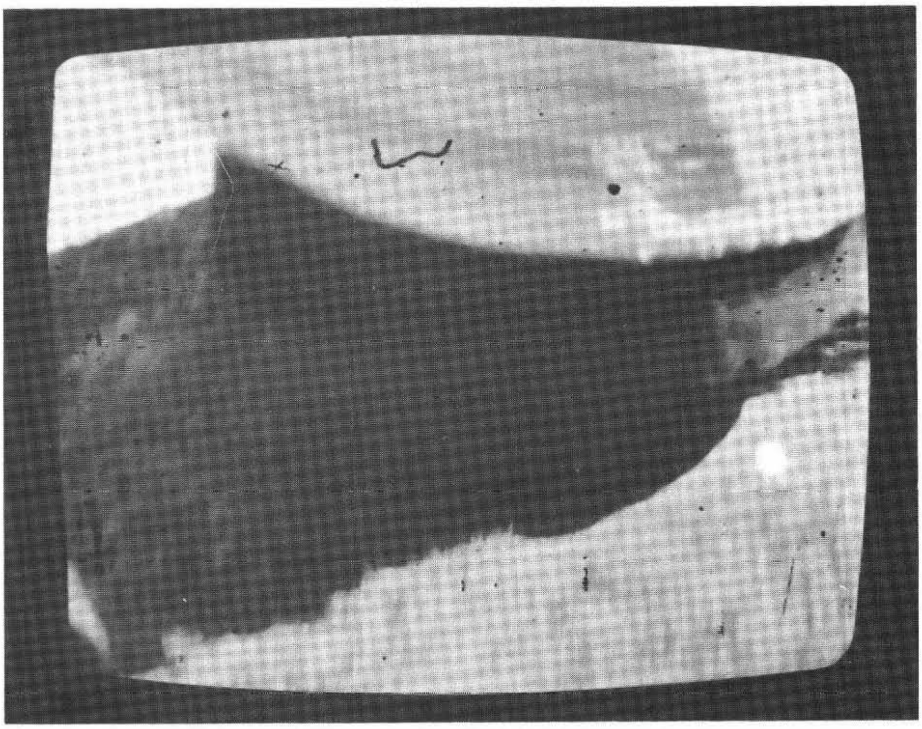

Fig. 28 – Errors caused by previous-frame substitution in the presence of motion

(a) dirt concealment inhibited

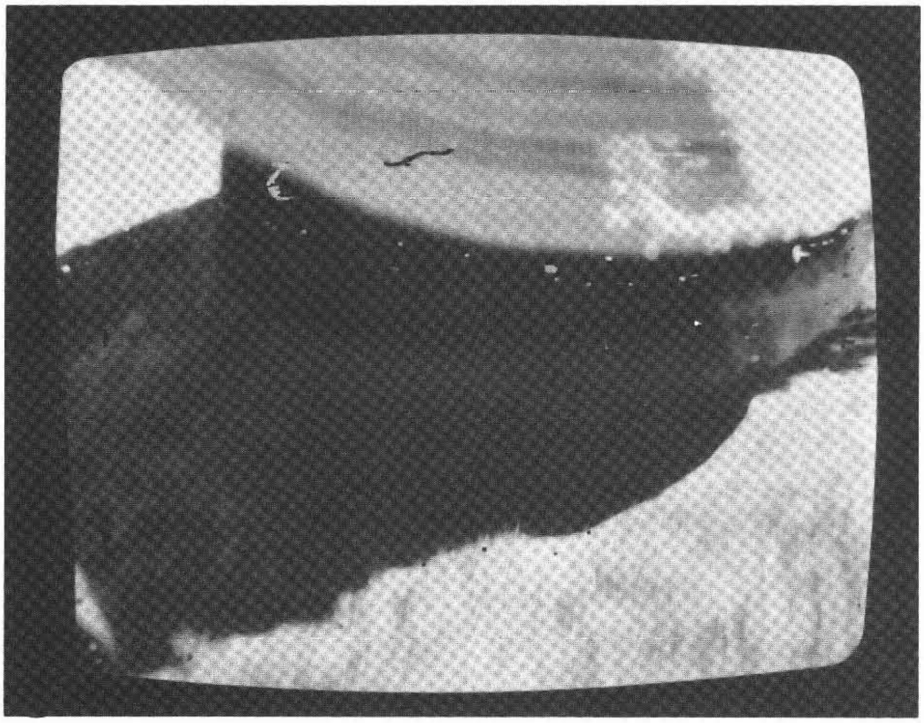

(b) dirt concealment enabled. Substitution errors are clearly visible near the light/dark boundary in the upper part of the frame

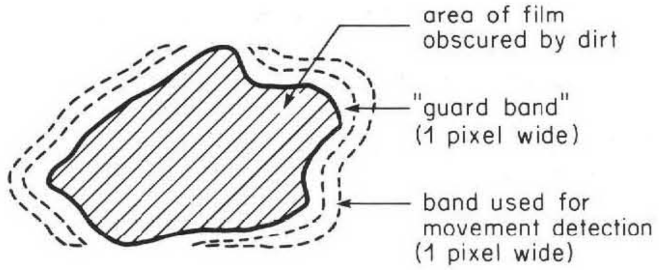
area of film obscured by dirt
"guard band" (1 pixel wide)
band used for movement detection (1 pixel wide)

Fig. 29 – Area of film used for movement detection.

Fig. 30(a) – Suppression of previous-frame substitution errors using the output from a movement detector.

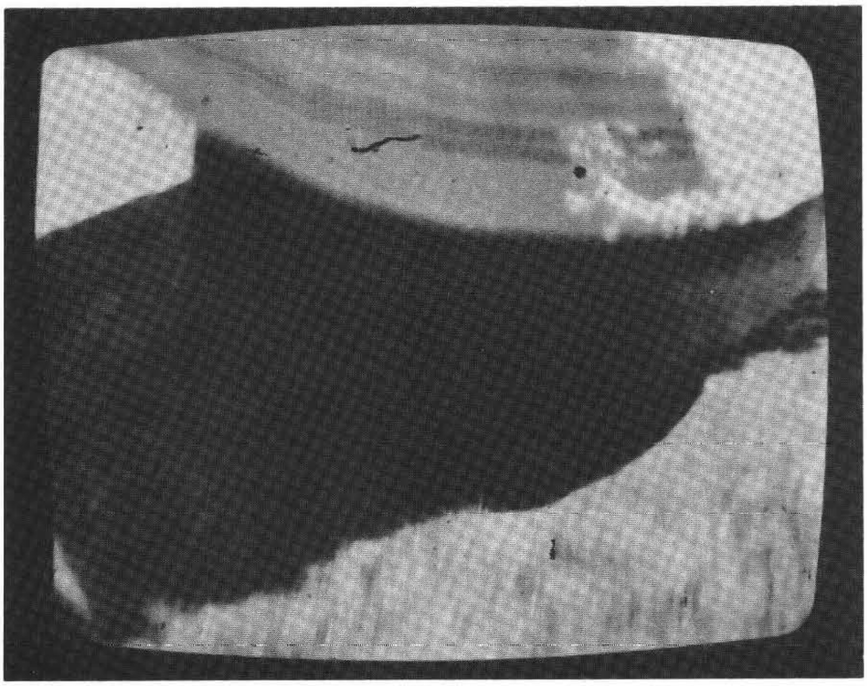

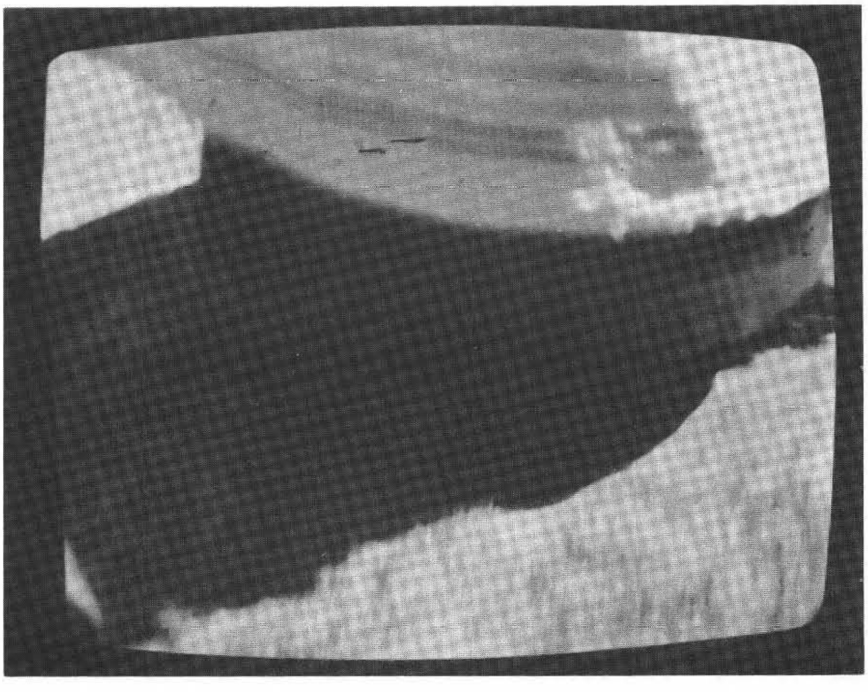

Fig. 30(b) – Output of two-stage dirt concealment system.

impression from a still picture as shown in Fig. 30(a); in practice, on moving material, the reappearance of dirt was judged much less objectionable than the presence of the artefacts shown in Fig. 28(b).

As a final stage in the investigation of concealment algorithms, the system was modified so that the motion detector, instead of simply inhibiting concealment completely, was used to select an alternative concealment algorithm. This "back-up" algorithm replaced corrupted areas of the picture with information interpolated from uncorrupted areas immediately adjacent to the blemish (see Fig. 31). Although it was known (ref. 14) that such a concealment algorithm was not always suitable due to its poor performance in detailed areas of a picture, it was nevertheless found that it was an adequate choice for a "back-up" algorithm since it was only used in those instances when the primary algorithm would have failed. Fig. 30(b) shows the performance of the complete two-stage dirt concealment system on the same film frame as Figs. 28 and 30(a), and it may be seen that the final result achieved is a considerable improvement on the original blemished film (Fig. 28(a)).

5. Conclusions

In recent years there have been several significant advances in the use of solid-state line-arrays for film scanning. As well as improvements in the sensitivity, spectral response and dynamic range of the sensors themselves, techniques such as correlated double sampling have allowed the performance of this type of telecine to be considerably improved. In addition, more sophisticated methods have been used to correct the output of the sensors for the effects of element-to-element sensitivity variations and optical shading.

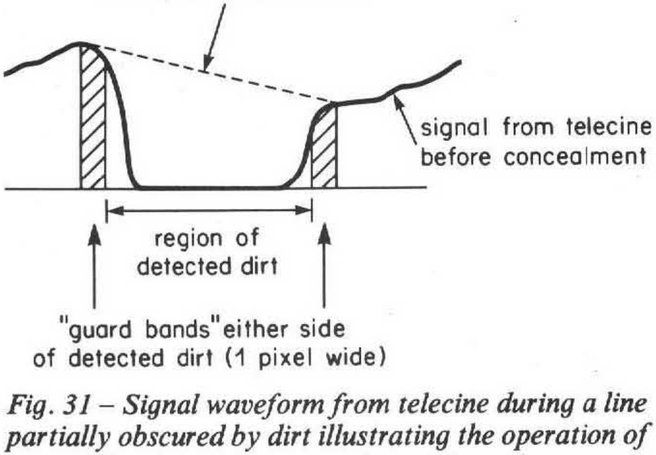
interpolated signal used to replace corrupted information
signal from telecine before concealment
region of detected dirt
"guard bands"either side of detected dirt (1 pixel wide)

Fig. 31 – Signal waveform from telecine during a line partially obscured by dirt illustrating the operation of the alternative concealment algorithm.

At the same time, the range of facilities available from a solid-state telecine has been expanded. Variable-speed operation over an extremely wide range of film speeds has been shown to be possible. And, finally, the good infra-red sensitivity of line arrays has been used to detect the presence of dirt and scratches on colour film and to enable suitable substitution methods to be used to conceal them.

6. Acknowledgements

The author would like to acknowledge the considerable contributions made by P. Mance, J.H. Moore and P.D. Ransome to the work described in this Report.

The work was carried out as part of a collaborative agreement with Rank Cintel Limited and many of the facilities described in this Report are included in the Cintel ADS-1 line-array telecine, launched at IBC '82[15]. A photograph of the ADS-1 telecine is shown in Fig. 32.

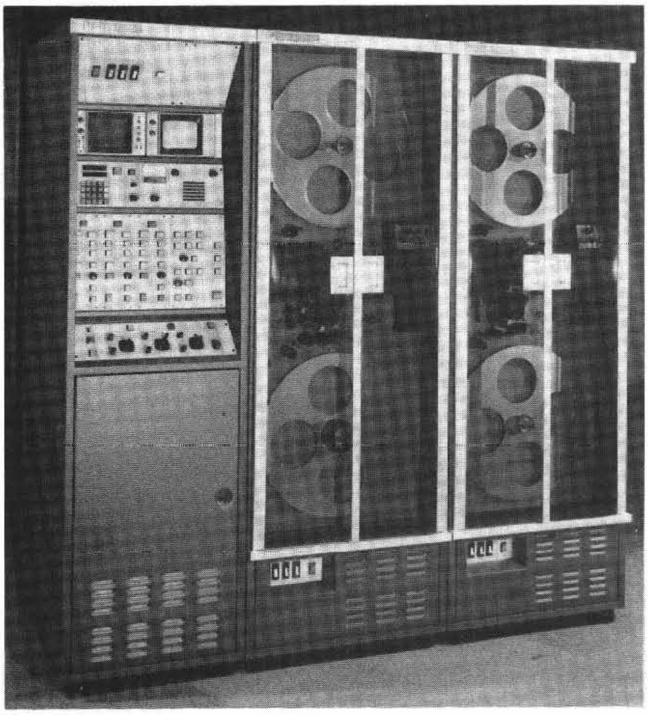

Fig. 32 – The ADS-1 line-array telecine developed by Rank Cintel Limited in collaboration with the BBC.

7. References

1. WRIGHT, D.T., 1973. Solid-state sensors: the use of a single dimension 512 element array for film scanning. BBC Research Department Report1973/32.

2. LE COUTEUR, G.M., 1976. Solid-state image sensors: improvements in signal processing techniques. BBC Research Department Report 1976/4.

3. LE COUTEUR, G.M., 1976. Performance tests on a selection of 100-element charge-coupled devices. BBC Research Department Report 1976/23.

4. CHILDS, I. *and* SANDERS, J.R., 1978. An experimental telecine using a line-array CCD sensor. SMPTE Journal Vol. 87, No. 4, pp. 209–213.

5. CHILDS, I., 1980. Colour operation of a solid-state sensor telecine. BBC Research Department Report 1980/11.

6. SEQUIN, C.H. *and* TOMPSETT, M.F., 1975. Charge transfer devices. Supplement 8 to Advances in Electronics and Electron Physics, New York, Academic Press, 1975.

7. KOSONOCKY, W.F. *and* CARNES, J.E., 1972. Two-phase charge-coupled shift registers. 1972 IEEE International Solid State Circuits Conference Digest, p. 132, IEEE, New York.

8. KRAMBECK, R.H., WALDEN, R.H. *and* PICKAR, K.A., 1971. Implanted barrier two-phase charge-coupled device. Applied Phys. Letts. 1971, Vol. 19, No. 12, pp. 520–522.

9. WALDEN, R.H. *et al.*, 1972. The buried-channel charge-coupled device. Bell Syst. Tech. J., Vol. 51, No. 7, pp. 1634–1640.

10. ESSER, L.J.M., 1973. The peristalitic charge-coupled device. Proceedings of the Charge-Coupled Device Applications Conference, San Diego, 1973, pp. 269–177.

11. OLIPHANT, A., 1978. A digital telecine processing channel. BBC Research Department Report 1978/9.

12. JONES, P.J., private communication.

13. LE COUTEUR, G.M., 1976. Improvements relating to telecine machines. British Patent No. 154 7811.

14. LE COUTEUR, G.M., 1976. Improvements relating to telecine machines. British Patent No. 154 7812.

15. CHILDS, I. *and* GRIFFITHS, M.J., 1982. Novel uses of digital processing in a modern telecine. Proceedings of the 1982 International Broadcasting Convention, Brighton, England, pp. 46–50.

Printed by BBC RESEARCH DEPARTMENT, Kingswood Warren, Tadworth, Surrey, KT20 6NP